

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 671

April, 1982

MULTI-LEVEL RECONSTRUCTION OF VISUAL SURFACES:

**Variational Principles and
Finite Element Representations**

Demetri Terzopoulos

Abstract

Computational modules early in the human vision system typically generate sparse information about the shapes of visible surfaces in the scene. Moreover, visual processes such as stereopsis can provide such information at a number of levels spanning a range of resolutions. In this paper, we extend this multi-level structure to encompass the subsequent task of reconstructing full surface descriptions from the sparse information. The mathematical development proceeds in three steps. First, the surface most consistent with the sparse constraints is characterized as the solution to an optimal approximation problem which is posed as a variational principle describing the constrained equilibrium state of a thin flexible plate. Second, local, finite element representations of surfaces are introduced and, by applying the finite element method, the continuous variational principle is transformed into a discrete problem in the form of a large system of linear algebraic equations whose solution is computable by local-support, cooperative mechanisms. Third, to exploit the information available at each level of resolution, a hierarchy of discrete problems is formulated and a highly efficient multi-level algorithm, involving both intra-level relaxation processes and bi-directional inter-level local interpolation processes is applied to their simultaneous solution. Examples of the generation of hierarchies of surface representations from stereo constraints are given. Finally, the basic surface approximation problem is revisited in a broader mathematical context whose implications are of relevance to vision.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505 and in part by National Science Foundation Grant 79-23110MCS. The author was supported by postgraduate scholarships from the Natural Sciences and Engineering Research Council of Canada and from Fonds F.C.A.C. pour l'aide et le soutien à la recherche, Québec, Canada.

1. INTRODUCTION

A fundamental problem in early vision is that of inferring the three-dimensional geometry of visible surfaces in a scene from the intensity information available in the retinal images. It seems that advanced biological vision systems subdivide this surprisingly difficult problem, distributing its solution over a number of subsystems or modules which contribute to the recovery of information about surface shape. Each module specializes in the interpretation of specific classes of image cues and, to a first approximation, it performs its task independently of the other subsystems. Examples of primary modules which have been identified include those responsible for stereo vision and the perception of motion.

The computational framework set forth by Marr [Marr, 1976, 1982; Marr and Poggio, 1977] has had a strong influence in the understanding of early vision. In particular, it has provided a paradigm which dictates that the modules initially be characterized in terms of the visual tasks which they must perform and, subsequently, that they be studied in terms of the computational processes through which they perform these tasks. The various computational processes in early vision transform symbolic representations of images into symbolic representations of surfaces over several stages of analysis. An understanding of some of the details of these processes and representations has evolved recently, especially at the stages closer to the image, where much of the work has been inspired by recent advances in neuroscience ([Marr, 1976], [Marr and Hildreth, 1980], [Marr and Poggio, 1979], [Ullman 1979a], etc.). On the other hand, our insights at stages closer to explicit surface and volumetric representations are meager. For example, these fundamental questions remain open: first, how is the information generated by the various modules, amalgamated into representations of surfaces (see, e.g., [Nishihara, 1981]) and, second, how do such representations give rise in turn to representations of the three-dimensional properties of objects in the scene (see, e.g., [Brooks, 1981], [Marr and Nishihara, 1978], and [Nevatia and Binford, 1977]).

The work described in this paper is part of ongoing research into the problem of obtaining surface representations which will be of use to later processing stages in vision. In the context of Marr's framework, our goal is to analyze the process through which the sparse information retrieved by, say, stereopsis or analysis of motion is combined and transformed into full, retinocentric descriptions of surface shape consistent with our perception when we look around us. In particular, *we will argue that multiple, full surface representations spanning a range of resolutions are desirable and, indeed, show that they may be generated as an integral part of a highly-efficient, multi-level surface reconstruction algorithm.* Moreover, our approach seems sufficiently general to allow several classes of surface shape information (such as local depth or orientation) provided by a number of vision modules to be merged in a meaningful way.

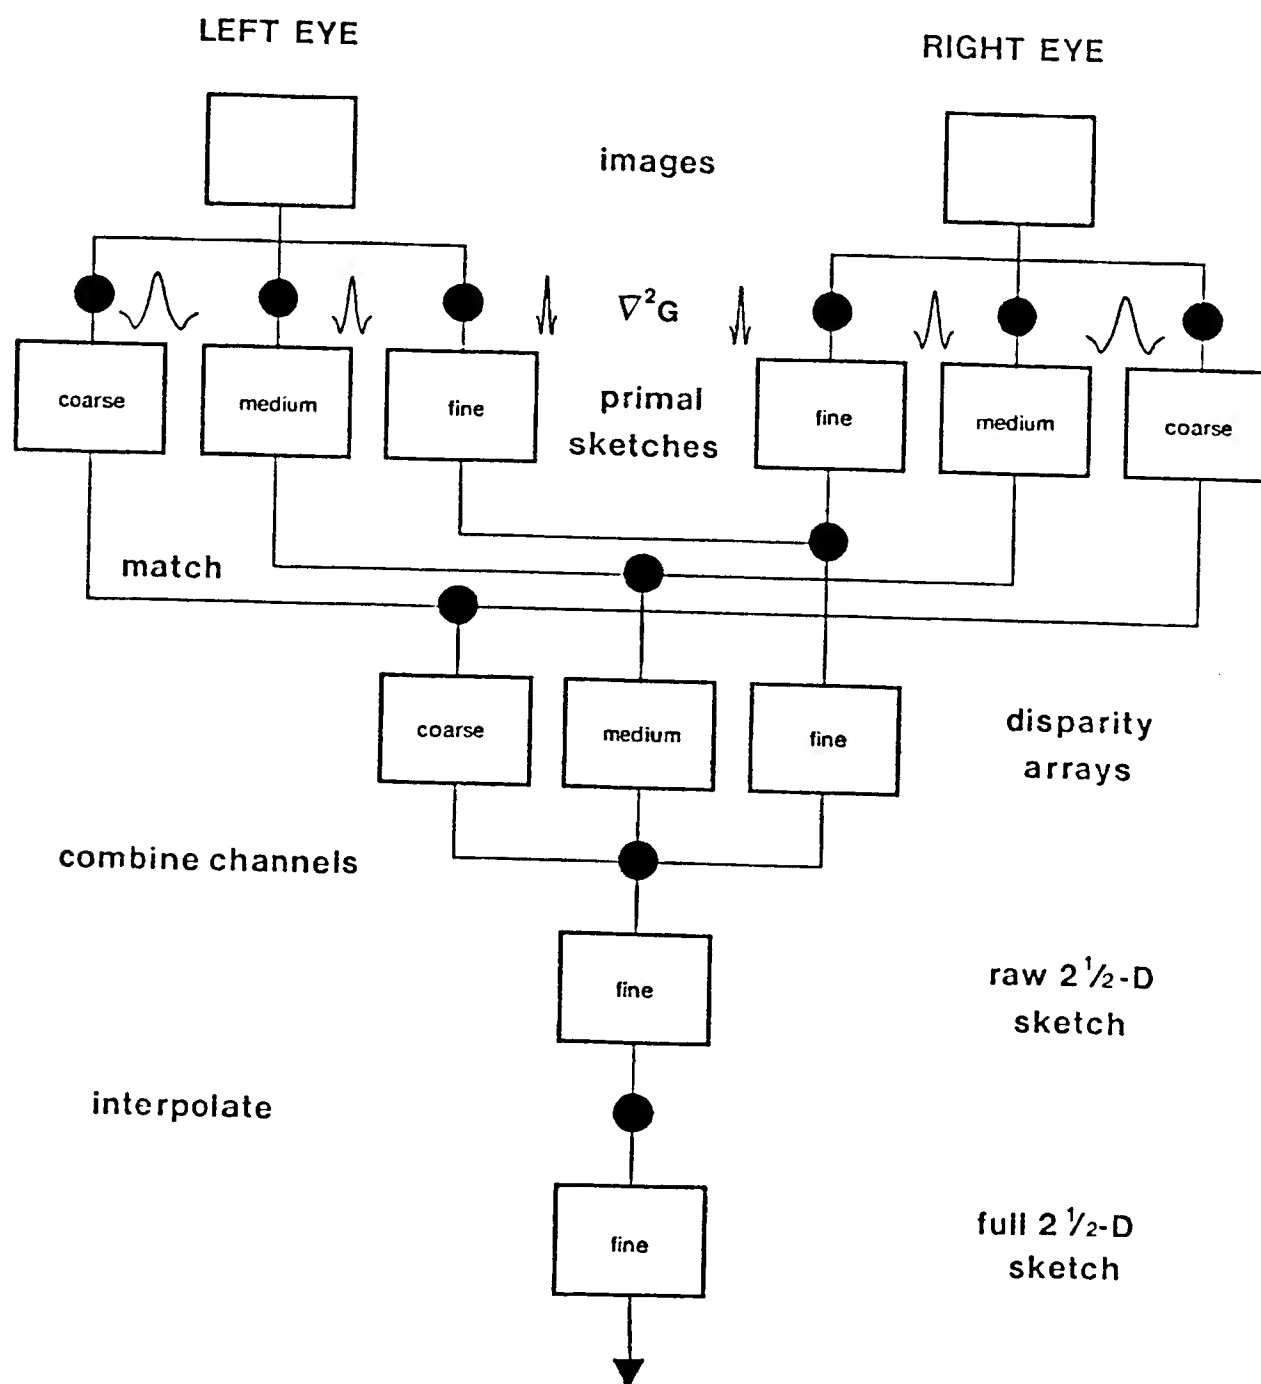
1.1. Motivation of the Multi-Level Approach

To clarify our intentions, we will first examine Marr's framework for early vision in some detail. The framework is characterized by at least three major processing stages, each of which transforms one retinocentric representation into another with the purpose of inferring, and making explicit, relevant information about the surfaces in a scene. The first stage transforms the intensity representations (or retinal images) into a primary representation, called the *primal sketch* [Marr, 1976]. Changes in the physical properties of surfaces almost always give rise to intensity changes in the images, and it is at the level of the primal sketch that the locations of these changes are made explicit. In the second processing stage, specialized processes, such as those concerned with stereo and shape from motion, infer information about the shape of surfaces from the contents of the primal sketch. Since inferences can typically be made only at those locations which have been marked in the primal sketch, the information generated is *sparse*, and it is collected into sparse representations of surface shape that are referred to as the *raw $2\frac{1}{2}$ -D sketch*. The final stage is one of *full surface reconstruction* in which the sparse representations are transformed into a *full $2\frac{1}{2}$ -D sketch* containing explicit information about surface shape at all points in the scene.

The goal of the first processing stage is the detection of intensity changes in the image. Recently, Marr and Hildreth [Marr and Hildreth, 1980; Hildreth, 1980] proposed a theory of edge detection which was inspired by existing neurophysiological evidence and certain mathematical issues. An important aspect of this theory is that intensity changes in the images must be isolated at different scales of resolution. Indeed, there is evidence that the human visual system detects intensity changes over a range of resolutions through the use of up to five independent, spatial-frequency-tuned, bandpass channels [Campbell and Robson, 1968; Wilson and Giese, 1977; Wilson and Bergen, 1979; Marr, Poggio, and Hildreth, 1980]. The existence of these independent primal sketch representations is a crucial factor which contributes to the success of some later computations such as stereopsis, as modeled by the Marr-Poggio theory of stereo vision [Marr and Poggio, 1979] (see also [Mayhew and Frisby, 1980, 1981]). According to this model, the bandpass nature of the channels leads to an almost trivial solution to the stereo correspondence problem within the disparity range of each channel. Detailed depth information over a wide disparity range is obtained through a process by which the coarser channels control vergence eye movements that bring the finer channels into alignment (general studies of vergence eye movements include [Riggs and Niehl, 1960] and [Rashbass and Westheimer, 1961]). On the other hand, computations such as motion correspondence [Ullman, 1979a], whose function may not depend critically on the existence of multiple representations, may nevertheless be operative at each of the levels. It seems likely in any case that multiple sparse representations of surface shape that span a range of resolutions are generated by most of these modules.

In the context of stereopsis, Grimson [1981a, 1982a] pioneered the mathematical theory of the subsequent visual surface reconstruction process which transforms the sparse surface

Figure 1. The stereo module with single-level surface reconstruction.



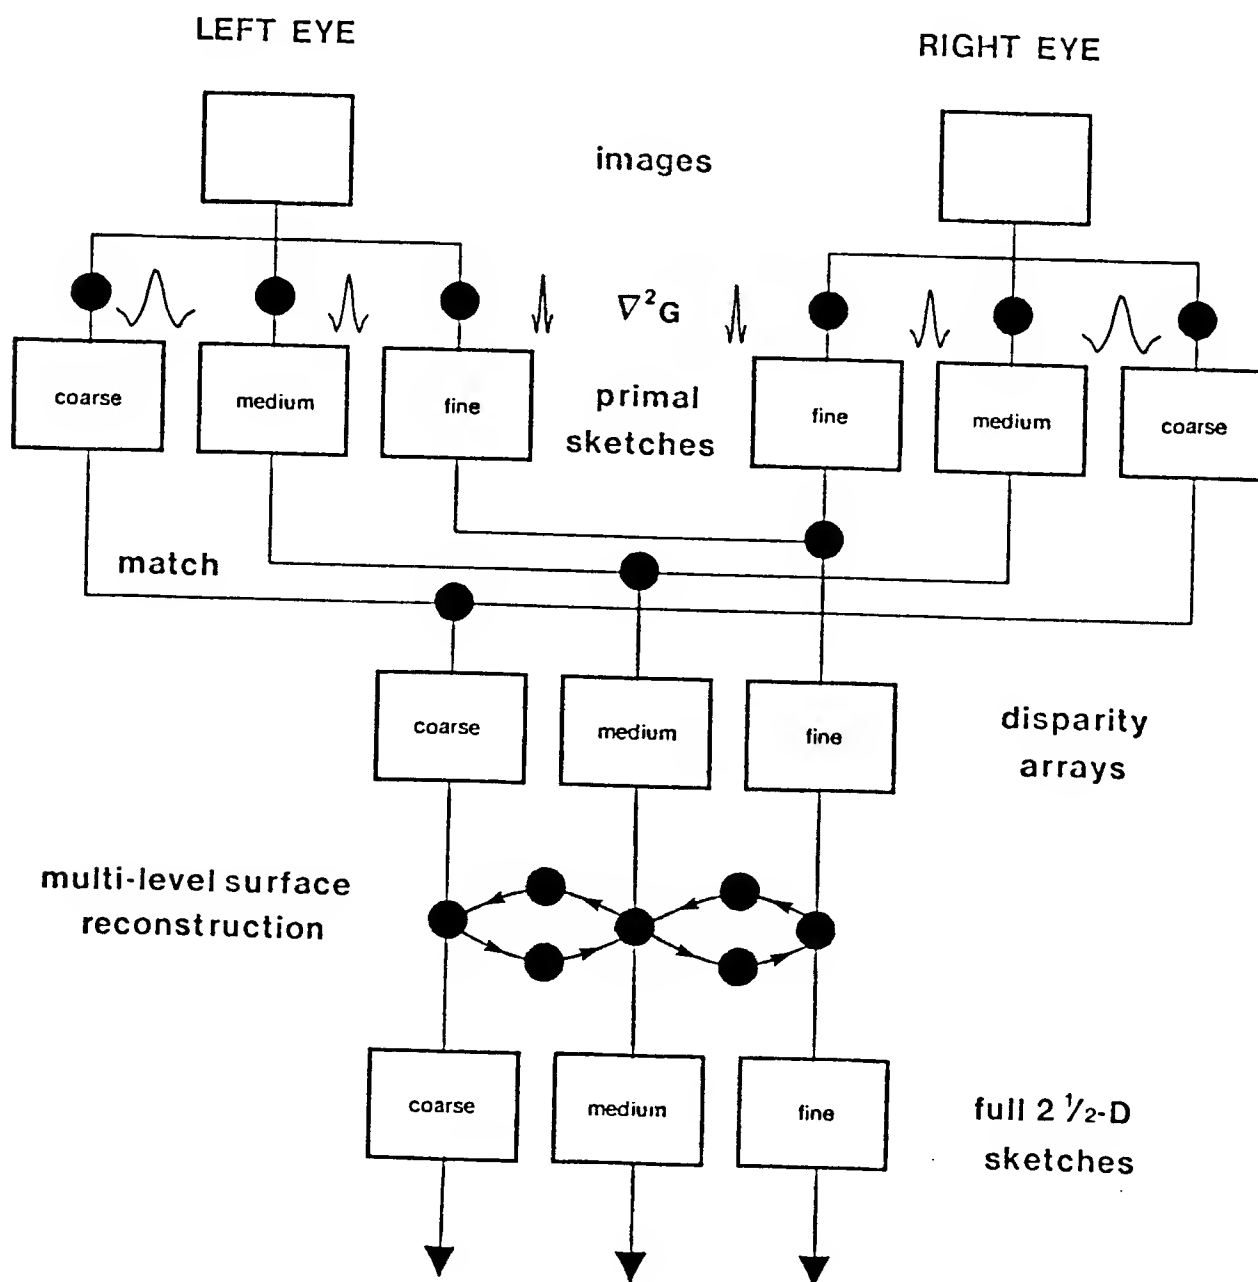
descriptions into full ones. He proposed that, before reconstruction begins, the multiple, sparse depth representations output through the different bandpass channels be combined into a single raw $2\frac{1}{2}$ -D sketch in a way which maintains consistency across all scales. The raw $2\frac{1}{2}$ -D sketch then contains sparse depth information at the finest resolution possible. Next, a single reconstruction process operating at this finest level generates a unique full $2\frac{1}{2}$ -D sketch representing depth information at high resolution. The steps are illustrated in Figure 1, in which only three bandpass channels are shown for simplicity.

A single full surface representation at the finest resolution possible certainly captures all of the information provided by the stereo module and it should, in principle, be sufficient input to later tasks. Unfortunately, a number of technical problems arise with this simple approach. First, in collapsing the multiple depth representations into one raw $2\frac{1}{2}$ -D sketch, information potentially useful in later processing stages which are concerned with object-centered surface descriptions, 3-D models of objects, and object recognition has been discarded prematurely. It now seems likely that in order for some of these later stages to succeed and work efficiently, surface representations at multiple scales will be necessary, just as they are necessary at earlier stages such as stereopsis. In accordance with Marr's *principle of least commitment* [Marr, 1976], it would be wasteful to discard information, prior to surface reconstruction, which may have to be regenerated later. A second, and more immediately serious problem is a consequence of the great bulk of incoming information within the large raw $2\frac{1}{2}$ -D sketch which must be processed at the finest resolution. Biologically feasible surface reconstruction algorithms such as those developed by Grimson are extremely inefficient at generating full surface descriptions when faced with such large representations. Roughly speaking, the primary reason for this inefficiency is due to the local nature of the algorithms in question.

The above problems may be avoided if the sparse representations are *not* collapsed into a single fine representation. Instead, multiple full surface representations spanning a range of resolutions ought to be generated by the reconstruction process itself and made available to processing stages beyond. The multi-level surface reconstruction algorithm which we will develop in this paper accomplishes precisely this. Because the algorithm exploits information available at coarser resolutions, its speed efficiency is dramatically superior to that of single level reconstruction schemes. Order-of-magnitude improvements are typically observed for surfaces reconstructed from information provided by stereopsis. On the other hand, the expense in space in maintaining all the coarser representations is very worthwhile since it turns out to be only a fraction of that required to maintain the finest one.

Figure 2 illustrates the multi-level surface reconstruction scheme and its incorporation into stereopsis. A fundamental point to realize about the multi-level approach in general is that information about surface depth, or for that matter surface orientation, is provided in each of the channels (i.e., sparse representations) by the various vision modules and, as will be shown, contributes in an optimal way to the generation of the hierarchy of full surface representations. The multi-level scheme involves both *intra-level* processes which propagate information across a single representation, as well as *inter-level processes* which communicate between representations. The inter-level processes are further classified into those which transfer information from coarser levels to finer ones, and those which transfer information from finer levels to coarser ones. At this point, we emphasize that multiple representations of consistent accuracy can be achieved only if such a *bi-directional* flow of information is allowed to take place between the levels. This statement will be substantiated rigorously in a later section.

Figure 2. Multi-level approach to surface reconstruction.



If the processes and representations envisioned are to be considered as models of the human visual system, their form is constrained from below by what can be implemented in neuronal hardware. Although our incomplete knowledge renders premature the formulation of precise arguments along these lines, some constraints which presently seem compelling, such as parallelism, locality and simplicity of computation, efficiency, uniformity, and extensibility [Ullman, 1979b] have been factors in the theoretical analysis of the surface reconstruction problem and in the search for algorithms (similar constraints are issues when considering implementations within parallel, pipelined computer architectures). Once a specific algorithm is selected based on these constraints, it may be implemented and its performance can be evaluated empirically

in terms of the original computational goal. If the performance is consistent with psychological evidence, the algorithm may be thought of as constituting an existence proof that the theory adequately models the human vision system. This is what we are ultimately striving for in our study of the surface reconstruction problem.

1.2. Overview

In this paper, we lay down the mathematical foundations of a multi-level approach to visual surface reconstruction, primarily in the context of stereo vision. With the help of a physical model, the basic surface approximation problem is given an intuitive interpretation. Although it is in general a nontrivial matter to solve this problem, our model suggests the application of potent methods which have arisen out of classical mathematical physics — the calculus of variations, optimal approximation theory, and functional analysis. Aspects of the above formalisms are employed to render our problem amenable to solution by numerical techniques. The development is as follows.

- In Chapter 2, visual surface reconstruction is cast as an optimal approximation problem which involves finding the equilibrium position of a thin flexible plate undergoing bending. The problem is posed formally as a variational principle which we propose to solve by first converting it to discrete form using the finite element method.
- In Chapter 3, we prepare the way for applying this discrete approximation method by finding a set of minimal conditions for our continuous problem to have a unique solution. We show that these requirements will almost always be satisfied in practice, so that we can consider our surface approximation problem to be well-posed, and can proceed to obtain the solution.
- In Chapter 4, we turn to the task of converting our continuous problem into discrete form. To do so, we define a simple nonconforming finite element which will constitute the basis of our local, piecewise continuous representation of surfaces. Because the element is nonconforming, we first must prove that it leads to unique discrete approximations, and that these approximations converge to the exact solution as the elements decrease in size. Having done this, we derive the discrete surface approximation problem as a large system of linear equations.
- In Chapter 5, we face the task of solving this linear system efficiently in a biologically-feasible way, and it is here that we motivate the multi-level approach for obtaining the solution. The approach involves setting up a hierarchy of discrete surface approximation problems which span a range of resolutions and exploit the information available at each scale, and subsequently invoking a multi-level algorithm to solve them simultaneously. We demonstrate the efficient performance of the multi-level surface reconstruction algorithm on constraints

from stereopsis, and demonstrate that it generates a useful hierarchy of accurate surface representations.

- In Chapter 6, we reexamine our surface reconstruction problem and show that it is a special case within a general class of optimal interpolation problems involving arbitrary degrees of continuity, in any number of dimensions. These general problems involve the minimization of functionals which possess a number of invariance properties making them attractive for application to problems in early vision whose solutions require the iterative propagation of smoothness constraints across retinocentric representations.
- In Chapter 7, we conclude by discussing the overall implications of our approach to issues concerning the isolation of depth discontinuities, and the incorporation of other sources of information such as surface orientation. We discuss possible solutions to these problems in view of our finite element representation of surfaces and the multi-level surface reconstruction algorithm.
- For convenience, in two of the appendices, we cover the relevant mathematical background of the finite element method and the iterative solution of large linear systems.

2. THE MOST CONSISTENT SURFACE

The sparse information about surface shape retrieved by the various vision modules is in general underconstraining. That is to say, it is insufficient to compel a unique inference of the physical properties of the surfaces in the scene. Yet, even when presented with impoverished stimuli (such as *random dot stereograms* [Julesz, 1971] or *kinetic depth effect displays* [Wallach and O'Connell, 1953; Wallach, 1959; Johansson, 1975; Ullman, 1979a; etc.]), the human visual system routinely arrives at unique interpretations and our typical perception is a stable one of full surfaces in depth. Clearly, the visual system must invoke certain assumptions which provide enough additional constraint to allow a unique full surface representation to be computed from the sparse information provided. However, these additional assumptions must be plausible in that they reflect certain domain-dependent expectations. For example, in stereo vision, the sparse information takes the form of depth constraints which embody measurements of the distances from the viewer to the surfaces of objects in the scene. The additional assumptions should then be based on general expectations about physical properties of surfaces in the visual world, as well as aspects of the optical and computational processes taking part in the generation of the depth constraints.

Grimson [1981a] explored a number of issues along these lines. Qualitatively, his thesis is as follows. A surface in the scene which varies radically in shape usually gives rise to intensity changes which are marked in the primal sketches as *zero-crossings* of the Laplacian of the Gaussian-convolved images ($\nabla^2 G * I$ — see [Marr and Hildreth, 1980; Hildreth, 1980]). Moreover,


it is only at the locations of zero-crossings that the Marr-Poggio stereo algorithm can generate measurements of the distance to the surface, in the form of explicit depth constraints. Therefore, the surface cannot in general be varying radically in depth between the constraints to which it gave rise. By introducing this additional surface smoothness assumption, the goal of accurately reconstructing the shape of visible surfaces and thereby computing full surface representations consistent with our perception is attainable in principle. A theoretical proof of this statement lies in the domain of mathematics. In the next section, we take the first step by rigorously formulating the surface reconstruction problem as an optimal approximation problem in which the smoothness assumption has a clear intuitive interpretation, and which eventually leads to efficient multi-level algorithms for its solution.

2.1. A Physical Interpretation — The Bending of a Thin Plate

Visual surface reconstruction can be characterized formally as a constrained, optimal approximation problem in two dimensions. In the context of stereo vision, where constraints embody depth measurements to surfaces in the scene, the goal is to reconstruct, as accurately as possible, the shape of the surface which gave rise to these measurements. Of course, it is necessary that we be able to deal with the complication arising from arbitrarily-placed constraints since, as has been argued, constraints of this type are generated naturally by stereopsis and other vision modules. More rigorously, the problem can be stated as follows: given a finite set of *arbitrarily located* distinct point constraints on a plane, each constraint having a real scalar value associated with it, find the unique *optimal* function of two variables which is *most consistent with these constraints*. Our notion of consistency will be defined shortly. We consider the solution to our problem to be a full surface representation in that it makes explicit our best estimate of the distance to every visible point on the surface in the scene.

The constraints provided by the stereo computation are never completely reliable. Errors due to noise, and errors in matching corresponding zero-crossings are bound to occur. This suggests that we should not try to interpolate the given data exactly because a few “bad” constraints can have a detrimental effect on the shape of the recovered surface. Relaxing the interpolation requirement turns our problem into one of surface approximation in which we would like to maintain control over how closely the surface fits the data.

By thinking in terms of an optimal surface, we imply the choice of a suitable criterion that will allow us to measure the optimality of admissible functions. A suitable criterion for measuring the optimality of surfaces in the context of surface approximation in stereopsis translates into a precise mathematical statement which captures intuitive notions about the smoothness of admissible surfaces as well as their closeness of fit to the known depth constraints. Perhaps the intuitively clearest treatment of our problem is in terms of a physical model. Consider a planar region Ω , the region within which we wish to obtain an optimal approximating surface most consistent with a finite set of sparse constraints. Let us imagine that the constraints constitute a

Figure 3. The physical model for surface approximation.


set of vertical pins scattered inside Ω , the height of an individual pin being related to the distance from the viewer to the surface in the scene. Suppose that we take a thin flexible plate of elastic material that is planar in the absence of external forces, and constrain it to pass near the tips of the pins by attaching ideal springs between the pin tips and the surface of the plate as shown in Figure 3. It is not difficult to imagine the equilibrium position of the plate as a function of the various pin heights.

Intuitively, the equilibrium position of the thin plate is a “fair” approximating surface in that it will exhibit a sufficient amount of smoothness between the constraints. Moreover, on quantitative grounds, there is evidence to suggest that such a surface is indeed an optimal one in terms of the imaging process [Grimson, 1982b]. In any case, we have a reasonable physical model for the optimal approximating surfaces and, moreover, this model will suggest good strategies for solving our problem.

We emphasize however that the appropriateness of the model depends on two important issues. The first involves ensuring that a unique solution exists, and the second is to guarantee that the solution is meaningful in view of the constraints. Firstly, we realize that the plate-spring system will be unstable for certain pin configurations. If we have but a single pin, then a stable equilibrium does not exist, as the plate has two unconstrained degrees of freedom (rotation about the axis of the pin is excluded). A similar degenerate situation arises for the case of any number of pins arranged linearly, the plate then having one unconstrained degree of freedom. Clearly

at least three noncollinearly arranged pins are required to assure that a unique state of stable equilibrium exists. Secondly, a reasonable choice must be made for the stiffness of the springs. If the springs are too strong in relation to the rigidity of the plate material, then a pin whose height deviates significantly from that of its neighbors (i.e. an erroneous constraint) will place an abnormally large amount of strain on the plate locally and have undesirable effects on the shape of the surface. On the other hand, if the springs are too weak, the intrinsic rigidity of the plate can overwhelm them and the plate will remain nearly planar over large variations in the height of the pins. In the limit of a rigid plate, the resulting planar, least-squares approximation would be meaningless in that the solution does not in general lie close to constraints other than those arising from nearly flat surfaces.

We will now proceed to a mathematical characterization of the above physical model. To do so, we apply the well-known *minimum potential energy principle* from classical mechanics, which states that the potential energy of a physical system in a state of stable equilibrium is at a local minimum. For the model, the potential energy in question is that due to deformation of the plate and springs, as well as the energy imparted by any externally applied forces. In this sense then, the surface we seek is one of minimal energy.

First, consider the plate. It is known (see, e.g., [Courant and Hilbert 1953, pg. 250], [Landau and Lifshitz, 1970]) that the potential energy of a thin plate under deformation is given by an integral of a quadratic form in the *principle curvatures* of the plate. If the principle curvatures of the deformed plate are denoted by κ_1 and κ_2 , the potential energy density is given by an expression of the form

$$\frac{A}{2}(\kappa_1^2 + \kappa_2^2) + B\kappa_1\kappa_2 = 2A\left(\frac{\kappa_1 + \kappa_2}{2}\right)^2 - (A - B)\kappa_1\kappa_2,$$

where A and B are constants determined by the plate material. The expression $\frac{1}{2}(\kappa_1 + \kappa_2)$ is the *first* or *mean curvature* and $\kappa_1\kappa_2$ is the *second* or *Gaussian curvature* of the plate's surface (see, e.g., [Hilbert and Cohn-Vossen, 1952]).

Let the function $v(x, y)$ denote the deflection of the plate normal to the region Ω which can be taken to lie in the X-Y plane. Assuming that the deflection function and its partial derivatives, v, v_x, v_y, \dots are small, it can be shown (see e.g. [Rektorys, 1969, pg. 368]) that

$$\frac{\kappa_1 + \kappa_2}{2} \approx \frac{1}{2}\Delta v, \quad \kappa_1\kappa_2 \approx v_{xx}v_{yy} - v_{xy}^2,$$

where $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ denotes the Laplacian operator. Thus, the potential energy density can be written in the following forms:

$$\begin{aligned}
e_1 &= \frac{1}{2}(\Delta v)^2 - (1 - \sigma)(v_{xx}v_{yy} - v_{xy}^2) \\
&= \frac{1}{2}(v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) + \sigma(v_{xx}v_{yy} - v_{xy}^2) \\
&= \frac{1}{2}\left[\sigma(\Delta v)^2 + (1 - \sigma)(v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2)\right],
\end{aligned} \tag{1}$$

apart from a multiplicative constant which depends on the physical properties of the elastic material of the plate, and which has been set to unity without loss of generality. The constant σ , called the *Poisson ratio*, measures the change in width as the material is stretched lengthwise.¹ The desired potential energy of deformation is obtained by integrating the energy density over the domain in question, and is given by

$$\mathcal{E}_1(v) = \int_{\Omega} \int_{\Omega} \frac{1}{2}(\Delta v)^2 - (1 - \sigma)(v_{xx}v_{yy} - v_{xy}^2) dx dy.$$

To the potential energy of deformation, we must add the potential energy due to any external forces which may be present. The energy due to a force density $g(x, y)$ applied to the surface of the plate (such as the effect of gravity) is given by

$$\mathcal{E}_2(v) = - \int_{\Omega} \int_{\Omega} gv dx dy.$$

External forces and bending moments may also be applied around the boundary $\partial\Omega$ of Ω . The energy due to a force density $p(s)$ on the boundary (s denotes arc length along the boundary) is

$$\mathcal{E}_3(v) = - \int_{\partial\Omega} p(s)v ds,$$

while the energy due to bending moments applied around the boundary is

$$\mathcal{E}_4(v) = - \int_{\partial\Omega} m(s) \frac{\partial v}{\partial n} ds,$$

where $m(s)$ is the density of applied bending moments normal to the curve and $\frac{\partial}{\partial n}$ denotes the directional derivative along the outward normal to $\partial\Omega$.

Finally, we must account for the potential energy of deformation of the springs. Let C denote the set of points in Ω at which the imaginary pins are located; that is, the sparse set of locations at which the surface is constrained. Furthermore, denote the height of the pin (the value of the constraint) by real scalars $c_{(x_i, y_i)}$ and the stiffness of the spring attached to it (influence of the constraint) by positive constants $\beta_{(x_i, y_i)}$ for all $(x_i, y_i) \in C$. According to Hooke's law for an ideal spring, the total potential energy of deformation in the springs is given by

¹From the last expression in (1), it is apparent that the potential energy density may be considered to be a *convex combination* with parameter σ of the square of the Laplacian and a quadratic term in second-order partial derivatives of the deflection function, $(v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2)$. This fact will be used in a subsequent discussion.

$$\mathcal{E}_5(v) = \frac{1}{2} \sum_{(x_i, y_i) \in C} \beta_{(x_i, y_i)} [v(x_i, y_i) - c_{(x_i, y_i)}]^2.$$

The equilibrium state of the mechanical system can be obtained as the solution to the following minimization problem which is referred to as a *variational principle*:

The deflection of the plate at equilibrium is that function u from a set V of admissible functions v , for which the total potential energy

$$\mathcal{E}(v) = \mathcal{E}_1(v) + \mathcal{E}_2(v) + \mathcal{E}_3(v) + \mathcal{E}_4(v) + \mathcal{E}_5(v) \quad (2)$$

is minimal.

Thus, quantitatively, the "most consistent" surface which we seek is the one having minimal energy \mathcal{E} .

The visual surface approximation problem has been posed, in integral form, as a variational principle which is *quadratic* in that it involves terms that are at most quadratic in v and its derivatives. Through the formalism of the *calculus of variations* one can express the necessary condition for the minimum as the *Euler-Lagrange equation* which in this case can be shown to be a linear, self-adjoint, partial differential equation. Much of classical mathematical physics is based on this duality, and it provides numerous techniques for solving our problem, those which are directed towards the variational principle, as well as those which are directed towards its Euler-Lagrange equation.

Whatever the strategy, although it is conceivable that the exact analytical solution u could be derived, it is normally impossible to do so for all but the simplest-shaped domains Ω . Consequently, we are led to consider a numerical approach in which we somehow convert our continuous problem into a discrete problem whose numerical solution closely approximates the exact continuous solution u . We propose to employ what is probably the most potent tool for obtaining discrete approximations currently available — the *finite element method*. The method is applied to the variational principle directly and, because the variational principle is quadratic, the resulting discrete problem will take the particularly simple form of a linear system of algebraic equations. The main advantage of the finite element method is its generality. In the context of our surface approximation problem, it can be applied over domains Ω of complicated shape, and it is *not* limited to uniform discretizations of these domains. The importance of the latter property in the context of vision is evident when one considers, for example, the nonuniform structure of the retina where it is known that resolution decreases approximately linearly with eccentricity (see [Wilson and Giese, 1977] and [Wilson and Bergen, 1979] for a quantitative model of this phenomenon in terms of the spatial-frequency channels in early vision). Moreover, the finite element method leads to linear systems which are readily solvable in a parallel, iterative fashion by a sparsely-interconnected network of simple processors, a mechanism which seems prevalent in early vision.

For several reasons, we have avoided the alternate route of using the well-known *finite difference method* to discretize the associated Euler-Lagrange equations (see, e.g., [Collatz, 1966], [Forsythe and Wasow, 1960], [Smith, 1977]). The finite difference method is much more restrictive in that it practically limits us to uniform discretizations, the underlying convergence theory is much less well developed and, perhaps most importantly, it becomes very difficult to discretize the natural boundary conditions associated with our surface approximation problem, a task which is done trivially by the finite element method.

The mathematical background of the finite element method that is of relevance to our problem is included in Appendix A for convenience. The appendix introduces the required theory and lists the fundamental theorems which we will invoke in applying the method to the task at hand. The process will consist of several steps. First we pose the variational principle in an abstract form that is the basis of the mathematical machinery presented in Appendix A. Next, we determine formally the requirements on the boundary conditions that must be satisfied to ensure that our variational principle is *well-posed*; i.e., has a unique solution. Only then can we proceed to apply the finite element method to approximate the solution.

3. THE VARIATIONAL PRINCIPLE

In this chapter we analyze the continuous variational principle which embodies our visual surface approximation problem, in preparation for the application of the finite element method. In view of the formalism presented in Appendix A, our first goal is to state the variational principle in the abstract form; that is, to isolate the energy inner product which characterizes our minimization problem. We then derive the associated Euler-Lagrange equation and, in the process, consider the various forms of boundary conditions that can be imposed. Finally, we choose the appropriate form of these conditions in view of our visual surface approximation problem and obtain formally the minimum requirements for our variational principle to be well-posed.

3.1. The Energy Inner Product

According to equation (2.2), our variational principle asserting that the equilibrium state of the thin plate is a minimal-energy configuration, may be stated mathematically as the minimization of the expression

$$\begin{aligned} \mathcal{E}(v) = & \int_{\Omega} \int_{\Omega} \frac{1}{2} (\Delta v)^2 - (1 - \sigma)(v_{xx}v_{yy} - v_{xy}^2) - gv \, dx \, dy \\ & - \int_{\partial\Omega} p(s)v \, ds - \int_{\partial\Omega} m(s) \frac{\partial v}{\partial n} \, ds + \frac{\beta}{2} \sum_{(x_i, y_i) \in C} [v(x_i, y_i) - c(x_i, y_i)]^2. \end{aligned} \quad (1)$$

Here, we have assumed that the spring stiffnesses $\beta_{(x_i, y_i)} = \beta$ for all $(x_i, y_i) \in C$. The admissible space V for our variational principle is in general a subspace of the second-order Sobolev space $H^2(\Omega)$ over the region Ω (refer to the discussion in Appendix A). If $u \in V$ minimizes \mathcal{E} , then

$\mathcal{E}(u) \leq \mathcal{E}(u + \epsilon v)$ for any $v \in V$. Therefore, to obtain the necessary condition for the minimum, we substitute for v the small variation $u + \epsilon v$ about u and equate to zero the derivative with respect to ϵ for $\epsilon = 0$. Equivalently, we may perform the variation using the rules of differentiation:

$$\begin{aligned} \delta \mathcal{E}(u) = & \int \int_{\Omega} \Delta u \delta(\Delta u) - (1 - \sigma)[u_{xx} \delta(u_{yy}) + u_{yy} \delta(u_{xx}) - 2u_{xy} \delta(u_{xy})] - g \delta u \, dx \, dy \\ & - \int_{\partial\Omega} p(s) \delta u \, ds - \int_{\partial\Omega} m(s) \delta \left(\frac{\partial u}{\partial n} \right) ds + \beta \sum_{(x_i, y_i) \in C} [u(x_i, y_i) - c_{(x_i, y_i)}] \delta u(x_i, y_i). \end{aligned}$$

Since variation commutes with differentiation, $\delta(\Delta u) = \Delta(\delta u)$, $\delta\left(\frac{\partial u}{\partial n}\right) = \frac{\partial}{\partial n}(\delta u)$, $\delta(u_{yy}) = (\delta u)_{yy}$, etc. If we now let $v = \delta u$, and set $\delta \mathcal{E} = 0$, we obtain

$$\begin{aligned} \int \int_{\Omega} \Delta u \Delta v - (1 - \sigma)(u_{xx} v_{yy} + u_{yy} v_{xx} - 2u_{xy} v_{xy}) - g v \, dx \, dy \\ - \int_{\partial\Omega} p(s) v \, ds - \int_{\partial\Omega} m(s) \frac{\partial v}{\partial n} ds + \beta \sum_{(x_i, y_i) \in C} [u(x_i, y_i) - c_{(x_i, y_i)}] v(x_i, y_i) = 0. \end{aligned} \quad (2)$$

Equations (1) and (2) may be cast in our abstract variational formulation of Appendix A. The key is in identifying the energy inner product as the bilinear form

$$a(u, v) = \int \int_{\Omega} \Delta u \Delta v - (1 - \sigma)(u_{xx} v_{yy} + u_{yy} v_{xx} - 2u_{xy} v_{xy}) \, dx \, dy + \beta \sum_{(x_i, y_i) \in C} u(x_i, y_i) v(x_i, y_i), \quad (3)$$

and in writing the linear form as

$$f(v) = \int \int_{\Omega} g v \, dx \, dy + \int_{\partial\Omega} p(s) v \, ds + \int_{\partial\Omega} m(s) \frac{\partial v}{\partial n} ds + \beta \sum_{(x_i, y_i) \in C} \left[c_{(x_i, y_i)} v(x_i, y_i) - \frac{1}{2} c_{(x_i, y_i)}^2 \right]. \quad (4)$$

Clearly then, (1) asserts that we are to minimize the quadratic functional $\mathcal{E}(v) = \frac{1}{2} a(v, v) - f(v)$, as required in the definition of the abstract variational principle (Definition A.1). On the other hand, (2) which expresses the necessary condition for the vanishing of the first variation may be written as $a(u, v) = f(v)$, as expected from the discussion leading up to the variational equation (A.10).

Having obtained expressions for the bilinear and linear forms, we can proceed to bring the finite element method to bear on the problem. Before doing so, however, it is imperative that we carry the analysis further so that we can express the necessary condition for a minimum as a partial differential equation, explore the issue of boundary conditions, and ensure that the problem is well-posed.

3.2. The Euler-Lagrange Equation and Boundary Conditions

For the duration of this section, we will ignore the summation term arising from the (spring) constraints, since its presence will complicate the notation while making no significant contribution to the discussion. First, we will transform the energy inner product $a(\cdot, \cdot)$ given in (3) using

integration by parts in two dimensions; i.e. Green's theorem. Let \mathbf{n} be the outward normal vector to $\partial\Omega$, \mathbf{t} be the usual tangent vector along $\partial\Omega$, and $\frac{\partial}{\partial n}$ and $\frac{\partial}{\partial t}$ denote partial differentiation along the normal and tangent respectively. Assuming that u is fourth-order differentiable, Green's identity (see, e.g., [Ciarlet, 1978, pg. 14], [Rektorys, 1969])

$$\int \int_{\Omega} u \Delta v - v \Delta u \, dx \, dy = \int_{\partial\Omega} u \frac{\partial v}{\partial n} - v \frac{\partial u}{\partial n} \, ds$$

may be used to transform the term $\Delta u \Delta v$ arising from the mean curvature of the surface:

$$\int \int_{\Omega} \Delta u \Delta v \, dx \, dy = \int \int_{\Omega} v \Delta^2 u \, dx \, dy + \int_{\partial\Omega} \Delta u \frac{\partial v}{\partial n} \, ds - \int_{\partial\Omega} v \frac{\partial}{\partial n} (\Delta u) \, ds, \quad (5)$$

where $\Delta^2 u = \Delta \Delta u = u_{xxxx} + 2u_{xxyy} + u_{yyyy}$. On the other hand, the Gaussian curvature term can be transformed using the identity [Ciarlet, 1978, pg. 15]

$$\int \int_{\Omega} u_{xx} v_{yy} + u_{yy} v_{xx} - 2u_{xy} v_{xy} \, dx \, dy = \int_{\partial\Omega} \frac{\partial^2 u}{\partial t^2} \frac{\partial v}{\partial n} \, ds - \int_{\partial\Omega} \frac{\partial^2 u}{\partial n \partial t} \frac{\partial v}{\partial t} \, ds. \quad (6)$$

If the boundary is sufficiently smooth, it can be shown (see [Rektorys, 1980, pp. 268–269]) that the second boundary integral can be written as

$$\int_{\partial\Omega} \frac{\partial^2 u}{\partial n \partial t} \frac{\partial v}{\partial t} \, ds = \int_{\partial\Omega} \frac{d}{ds} \left(\frac{\partial^2 u}{\partial n \partial t} \right) v \, ds. \quad (7)$$

Substituting equations (5)–(7) into (3) (and ignoring the constraint term), we obtain

$$a(u, v) = \int \int_{\Omega} v \Delta^2 u \, dx \, dy + \int_{\partial\Omega} P(u) v \, ds + \int_{\partial\Omega} M(u) \frac{\partial v}{\partial n} \, ds,$$

where

$$P(u) = -\frac{\partial}{\partial n} (\Delta u) + (1 - \sigma) \frac{d}{ds} \left(\frac{\partial^2 u}{\partial n \partial t} \right);$$

$$M(u) = \Delta u - (1 - \sigma) \frac{\partial^2 u}{\partial t^2}.$$

Thus, the necessary condition for the minimum (2) becomes

$$\int \int_{\Omega} (\Delta^2 u - g) v \, dx \, dy + \int_{\partial\Omega} [P(u) - p(s)] v \, ds + \int_{\partial\Omega} [M(u) - m(s)] \frac{\partial v}{\partial n} \, ds = 0$$

Now, since the above equation must hold and since v and $\frac{\partial v}{\partial n}$ are arbitrary on the closed region $\bar{\Omega}$, we must have

$$\Delta^2 u = g \quad \text{in } \Omega. \quad (8)$$

This is the fourth-order, linear, self-adjoint Euler-Lagrange equation that governs the small deflection of a thin plate at equilibrium, and it is satisfied by u inside Ω regardless of the boundary conditions on $\partial\Omega$. In its homogeneous form $\Delta^2 u = 0$, it is called the *biharmonic equation*. Furthermore, u must satisfy the *natural boundary conditions*

$$P(u) = p(s) \quad \text{and} \quad M(u) = m(s) \quad \text{on } \partial\Omega. \quad (9)$$

According to (6), the integral over Ω of the Gaussian curvature approximation $(v_{xx}v_{yy} - 2v_{xy}^2)$ has no effect whatsoever on the Euler-Lagrange equation, but contributes only to the boundary conditions.² This reflects the fact that the Gaussian curvature of a surface patch is invariant as the surface is subjected to arbitrary bending, and that its average value over the patch depends only on the tangent planes of the surface around the periphery of the patch [Hilbert and Cohn-Vossen, 1952, pp. 193-204]. This invariance property of the Gaussian curvature renders it inappropriate as a measure of surface consistency. For example, it cannot distinguish two developable surfaces such as a wildly undulating sinusoidal surface and a planar surface, both of which are cylinders and therefore have zero Gaussian curvature everywhere. On the other hand, the mean curvature does not in general remain invariant under bending and therefore plays a vital role in our energy inner product. This is evident from equation (1) — no value of σ can make $(\Delta v)^2$, which approximates the mean curvature, vanish.³

Another consequence of the necessary condition for the minimum is that the form of the natural boundary conditions satisfied by u are determined by any *essential boundary conditions* which may be imposed on v . In general, we can impose up to two essential boundary conditions, one on v and the other on $\frac{\partial v}{\partial n}$. First, consider the case of a *simply supported plate* where the essential boundary condition $v = 0$ is imposed on $\partial\Omega$ but $\frac{\partial v}{\partial n}$ is left unconstrained. The solution u must then still satisfy the second condition in (9). We therefore have the *Neumann* boundary conditions

$$u = 0, \quad M(u) = m(s) \quad \text{on } \partial\Omega,$$

and, moreover, the first contour integral in (1) vanishes.

Next, suppose that we also set $\frac{\partial v}{\partial n} = 0$ on $\partial\Omega$. Then,

$$u = \frac{\partial u}{\partial n} = 0 \quad \text{on } \partial\Omega, \quad (10)$$

which are the *Dirichlet* boundary conditions for the *clamped plate*. In this case, both contour integrals in (1) vanish and, moreover, σ is arbitrary since it does not appear in the Euler-Lagrange equation (8), but only in the natural boundary conditions (9) which have now been replaced by (10). We can therefore greatly simplify the variational integral. In particular, the functional minimization problems involving

²Expressions possessing this property are called *divergence expressions* [Courant and Hilbert, 1953].

³Brady and Horn [1981, pg. 29] state that "the choice of which performance index to use is reduced to the square Laplacian, the quadratic variation, and linear combinations of them". We stress that one should be careful not to choose that linear combination which results in a divergence expression (the Gaussian curvature) and therefore has an identically zero Euler-Lagrange equation. Recall from equation (2.1) that the small-deflection theory of the thin plate allows only a *convex combination* so, fortunately, it is free from danger in this respect.

$$\mathcal{E}(v) = \int \int_{\Omega} \frac{1}{2} (\Delta v)^2 - gv \, dx \, dy, \quad \text{for } \sigma = 1,$$

and

$$\mathcal{E}(v) = \int \int_{\Omega} \frac{1}{2} (v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) - gv \, dx \, dy, \quad \text{for } \sigma = 0,$$

are equivalent in the Dirichlet case.

Finally, consider the case of a free boundary; that is, when the externally imposed force $p(s)$ and moment $m(s)$ on the boundary are zero. Then, there are no constraints on v but, according to (9), u must satisfy

$$P(u) = M(u) = 0 \quad \text{on } \partial\Omega.$$

These are the natural boundary conditions satisfied by the solution for the case of the *free plate*. Once again, the contour integrals in (1) vanish and the energy functional takes the simple form

$$\mathcal{E}(v) = \int \int_{\Omega} \frac{1}{2} (\Delta v)^2 - (1 - \sigma)(v_{xx}v_{yy} - v_{xy}^2) - gv \, dx \, dy.$$

In general, the admissible space V is the subspace of the Sobolev space $\mathcal{H}^2(\Omega)$ which satisfies the essential boundary conditions imposed on the plate. If, for example, a portion of the edge of the plate is simply supported, V will consist of functions which satisfy the essential condition $v = 0$ on that portion of $\partial\Omega$. If part of the edge of the plate is clamped, then $v = \frac{\partial v}{\partial n} = 0$ on that part of $\partial\Omega$. On the other hand, if part of the edge is free, then no constraints are imposed on v over that portion of the boundary and, in the case of the free plate, $V \equiv \mathcal{H}^2(\Omega)$. Of course, the plate cannot be “too free” on $\bar{\Omega}$, because then the physical system cannot achieve stable equilibrium and a unique solution would not exist. Precisely how much freedom can be allowed will be established formally in the next section.

3.3. When is the Problem Well-Posed?

Turning to our visual surface approximation problem, we should at this point choose the appropriate form of boundary conditions on $\partial\Omega$. Since the only information about the surface that is provided by the stereo module, for example, is embodied in the sparse constraints, the strategy of least commitment is to “assume nothing” about the boundaries of the surface. In terms of our plate problem, this means that we should impose no essential boundary conditions on the plate; that is, we solve the *free plate problem* whose admissible space $V \equiv \mathcal{H}^2(\Omega)$.

If the boundary of the plate is free, it is clear that the constraints will play a critical role in providing a unique state of stable equilibrium. Our goal in this section is to specify the existence and uniqueness requirements mathematically as conditions under which the surface approximation problem is well-posed. To do this, we will invoke Theorem A.1, and satisfy its conditions by proving the following two propositions.

Proposition 1. The energy inner product $a(\cdot, \cdot)$ is symmetric.

Proof. $a(u, v) = a(v, u)$ is evident by inspection of equation (3). ■

Proposition 2. If the set of constraints C contains (at least) three noncollinear points, then $a(\cdot, \cdot)$ is V -elliptic for $0 \leq \sigma < 1$.

Proof. We want to show that there exists an $\alpha > 0$ such that $a(v, v) \geq \alpha \|v\|^2$, for all $v \in V$. To do so, it is sufficient to show that $a(v, v) = 0$ only if $v \equiv 0$. We rewrite $a(v, v)$ as

$$a(v, v) = \int_{\Omega} \sigma (\Delta v)^2 + (1 - \sigma) (v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) dx dy + \beta \sum_{(x_i, y_i) \in C} v(x_i, y_i)^2.$$

Now, $\Delta v = 0$ only if v is a *harmonic function*, while $(v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) = 0$ only if v is a first degree polynomial (as can easily be shown by integration), which is a subclass of the harmonic functions. Thus, the integral is ≥ 0 for $0 \leq \sigma < 1$ and it is zero only if v is a linear function over Ω . On the other hand, since β is positive by definition the sum is also ≥ 0 and it is zero only if $v(x_i, y_i) = 0$ for all $(x_i, y_i) \in C$. Therefore, if C contains three noncollinear points, then $a(v, v) = 0$ only if $v \equiv 0$, implying that $a(\cdot, \cdot)$ is V -elliptic. ■

By Propositions 1 and 2 and Theorem A.1, we are assured that the continuous approximation problem is well-posed if $0 \leq \sigma < 1$ and the set of constraints includes at least three noncollinear points. The condition on the constraints is not unexpected in view of the arguments made in Section 2.1. Physically speaking, all unconstrained degrees of freedom of the plate must be precluded, and three noncollinear constraints is clearly the minimum requirement for this to be the case. In application to natural images, the stereo algorithm will almost always generate at least three noncollinear points, so we can, for all practical purposes, consider our surface approximation problem to be well-posed so long as $0 \leq \sigma < 1$.

4. OBTAINING THE DISCRETE PROBLEM

So far, we have been dealing with the continuous form of our surface approximation problem. We formulated it in the required abstract form, selected appropriate boundary conditions, and showed that it is well-posed in practice. In this chapter, we face the task of applying the finite element method to transform the variational principle into an appropriate discrete problem whose discrete solution can be computed fairly easily. Our piecewise continuous representation of surfaces will be based on a very simple finite element which is, however, *nonconforming*. This will force us to introduce an approximate variational principle and to show that it has a unique solution which converges to the exact solution as the elements shrink in size. Only then can we undertake the next step which is to derive the discrete problem explicitly as a linear system of algebraic equations.

4.1. Conforming vs Nonconforming Methods

Our well-posed variational principle satisfies all the necessary conditions to guarantee that any *conforming* finite element method applied to it will converge. In principle, it is straightforward to apply a conforming finite element method according to the steps in Appendix A. We generate a finite element space S^h which is a *subspace* of our admissible space V , and apply the Ritz method to find that function $u^h \in S^h$ which optimally approximates the exact solution $u \in V$. The approximation is optimal in the sense that it is closest to u with respect to the strain energy norm $a(\cdot, \cdot)^{\frac{1}{2}}$, or equivalently, that the strain energy in the error $a(u - u^h, u - u^h)$ is minimal. To construct a conforming finite element subspace, we must satisfy the completeness and conformity conditions given in Section A.4. Since the energy inner product $a(\cdot, \cdot)$ contains partial derivatives of order $m = 2$, the completeness condition requires that the local polynomial defined within each element subdomain E must be at least a quadratic; $P^E \supset \Pi^2(E)$ for all $E \in \mathcal{T}^h$ ($\Pi^n(E)$ denotes the set of n^{th} degree polynomials over E). On the other hand, the conformity condition states that the polynomials must be of class C^1 across inter-element boundaries, and consequently $S^h \subset C^1(\bar{\Omega})$ globally. In satisfying both conditions, we are guaranteed that the finite element space is a subspace of the admissible space V , and that there exists a unique optimal approximation $u^h \in S^h$.

If $\bar{\Omega}$ is a polygonal region, elements with straight sides will suffice. A number of such elements which are conforming for $m = 2$ (i.e., problems characterized by fourth-order Euler-Lagrange equations) are available. Examples are the *Argyris triangle*, *Bell triangle*, and *Bogner-Fox-Schmidt rectangle* (see, e.g., [Ciarlet, 1978], [Strang and Fix, 1973], [Zienkiewicz, 1977] and the references therein). Unfortunately, we can expect serious computational difficulties to arise in the implementation of these conforming methods. The basic source of difficulty is the requirement of continuity of first partial derivatives across inter-element boundaries — either the structure of the conforming element spaces P^E becomes complicated, or their dimension is large. For our problem, the simplest conforming polynomial element is the Bell triangle, in which we have a quintic polynomial uniquely determined by 18 nodal variables consisting of the approximation v^h , as well as its first and second partial derivatives at the three vertices.

As is described in Appendix A, the dimensions of the finite element space can be reduced by the use of nonconforming elements. A popular nonconforming element for fourth-order problems is *Adini's rectangle*, whose local function p^E is a 12 degree-of-freedom polynomial with nodal variables being the approximating function, as well as its first partial derivatives at the four vertices. The element is nonconforming since it is only of class C^0 across inter-element boundaries. Many other nonconforming elements have been developed for fourth-order problems (see, e.g., [Ciarlet, 1978], [Strang and Fix, 1973], [Zienkiewicz, 1977]).

For this initial implementation, we have chosen to reduce the dimensions of the finite element space as much as possible by defining what for our problem is probably the simplest successful

nonconforming element imaginable. This element will be defined next.

4.2. A Simple Nonconforming Element

We will define a finite element space by the standard procedure outlined in Section A.4. Suppose that $\bar{\Omega}$ is rectangular, and consider a *uniform* triangulation \mathcal{T}^h of $\bar{\Omega}$ into identical square elements E , where the fundamental length h is the length of a side of E . By definition, we require that $\bigcup_{E \in \mathcal{T}^h} E = \bar{\Omega}$ and that the elements be adjacent and overlap along their sides. A point in $\bar{\Omega}$ is a *node* of the triangulation if it is a vertex of an elemental square and, as usual, we consider the elements to be inter-connected at the nodes. The nodal variables, will simply be the *node displacements*; i.e., the values of the function $v^h \in S^h$ at the nodes.

The next step is to define a space P^E of polynomials p^E over the element domain. The polynomials must satisfy the completeness condition which states that $\Pi^2 \subset P^E$, since the energy inner product contains derivatives of order $m = 2$. This is the requirement that the polynomials be able to reproduce exactly all states of constant strain which, in this case, are all polynomials up to degree two. We will satisfy this requirement by choosing P^E to be the six-dimensional space of full second degree polynomials $p^E: E \mapsto \mathbb{R}$ such that


$$p^E(x, y) = ax^2 + by^2 + cxy + dx + ey + f, \quad (1)$$

where the six real parameters a to f are to be determined.

We must ensure that p^E is uniquely specified within E in terms of the node displacements. To do so, we isolate a representative element and set the origin of the X-Y coordinate system at its lower left hand corner, as illustrated in Figure 4. Our task is to choose a p^E -*unisolvent* set of nodes, the displacements at which uniquely determine p^E . An appropriate choice is the six nodes shown in the figure, whose node displacements are denoted by $v_{i,j} \in \mathbb{R}$, for $(i, j) \in \{(-1, 0), (0, 0), (1, 0), (0, -1), (0, 1), (1, 1)\}$. Expressing the six unknown parameters in terms of the node displacements is then a simple matter of substituting the displacements into (1) and solving the resulting nonsingular system of six equations. We obtain

$$\begin{aligned} a &= \frac{1}{2h^2}(v_{1,0} - 2v_{0,0} + v_{-1,0}), \\ b &= \frac{1}{2h^2}(v_{0,1} - 2v_{0,0} + v_{0,-1}), \\ c &= \frac{1}{h^2}(v_{1,1} - v_{0,1} - v_{1,0} + v_{0,0}), \\ d &= \frac{1}{2h}(v_{1,0} - v_{-1,0}), \\ e &= \frac{1}{2h}(v_{0,1} - v_{0,-1}), \\ f &= v_{0,0}. \end{aligned} \quad (2)$$

Of course, the six degrees of freedom of this element are insufficient to enforce C^1 continuity of v^h across inter-element boundaries. Therefore, the element is nonconforming; $S^h \not\subset C^1(\Omega)$. It

Figure 4. Unisolvent nodes for the nonconforming element.


is a simple matter to show that the polynomials p^E are in general discontinuous across element boundaries, although continuity is maintained at the nodes themselves because each polynomial interpolates its unisolvent set of nodal displacements. At this point, we acknowledge that our element is somewhat unorthodox in that the definition of p^E requires nodal variables associated with two nodes which lie outside the element domain E . The justification for this transgression is that our element, as defined, will yield a discrete system whose matrix is particularly simple and uniform in structure. This will simplify the eventual implementation considerably. On the other hand, alternate arrangements for the unisolvent set of nodes are clearly possible. Perhaps a more appropriate choice from a biological standpoint would be a hexagonal triangulation with the unisolvent set of nodes placed at the vertices of hexagonal element domains E having sides of common length h . Regardless of our particular choice, the quadratic elements must first be shown to be convergent since they will invariably be nonconforming. This we will do in the next section for the simple square element.

4.3. The Approximate Variational Principle

Due to the nonconformity of the elements, $S^h \not\subset \mathcal{H}^2(\Omega) \equiv V$, and the finite element space is not a subspace of the admissible space V . Therefore, in lieu of the energy inner product $a(\cdot, \cdot)$ of equation (3.3), we are forced to substitute the approximate energy inner product $a_h(\cdot, \cdot)$, given by (A.16), which can be written as

$$\begin{aligned} a_h(u^h, v^h) &= \sum_{E \in \mathcal{T}^h} \int \int_E \Delta u^h \Delta v^h - (1 - \sigma)(u_{xx}^h v_{yy}^h + u_{yy}^h v_{xx}^h - 2u_{xy}^h v_{xy}^h) dx dy \\ &\quad + \beta \sum_{(x_i, y_i) \in C} u^h(x_i, y_i) v^h(x_i, y_i) \\ &= \sum_{E \in \mathcal{T}^h} \int \int_E \sigma \Delta u^h \Delta v^h + (1 - \sigma)(u_{xx}^h v_{xx}^h + 2u_{xy}^h v_{xy}^h + u_{yy}^h v_{yy}^h) dx dy \\ &\quad + \beta \sum_{(x_i, y_i) \in C} u^h(x_i, y_i) v^h(x_i, y_i), \end{aligned} \tag{3}$$

where $0 \leq \sigma < 1$. The corresponding approximate variational principle and variational equation are given by equations (A.17) and (A.18) respectively.

Does the approximate variational principle have a unique solution $u^h \in S^h$? To answer this question, we proceed in the spirit of Section A.5 by equipping S^h with a norm which we will employ to show that $a_h(\cdot, \cdot)$ is uniformly S^h -elliptic.

Proposition 3. The mapping $\|v^h\|_h: S^h \mapsto \mathbb{R}$ defined by

$$\|v^h\|_h = \left(\sum_{E \in \mathcal{T}^h} |v^h|_{2,E}^2 + \sum_{(x_i, y_i) \in C} v^h(x_i, y_i)^2 \right)^{\frac{1}{2}},$$

where $|v^h|_{2,E} = \left(\int \int_E (v_{xx}^h)^2 + (v_{xy}^h)^2 + (v_{yy}^h)^2 dx dy \right)^{\frac{1}{2}}$ is the second-order Sobolev semi-norm (see (A.3)), is a norm over S^h .

Proof. $\|\cdot\|$ is *a priori* only a semi-norm over S^h . Consider a $v^h \in S^h$ such that $\|v^h\|_h = 0$. Then it must be the case that (i) $|v^h|_{2,E} = 0$ for all $E \in \mathcal{T}^h$, and that (ii) $v^h(x_i, y_i) = 0$ for all $(x_i, y_i) \in C$. Because of their interpolatory nature, the local polynomials p^E are continuous at all the nodes. Moreover, by condition (i), v^h must be a first-degree polynomial inside every E . With $a = b = c = 0$ in equation (2), it is a simple matter to show that this implies that v^h is a continuous linear function over $\bar{\Omega}$. Now, by condition (ii), v^h is zero at all $(x_i, y_i) \in C$. Since the continuous problem is assumed to be well-posed, C contains at least three noncollinear points. Consequently $v^h \equiv 0$, and $\|\cdot\|_h$ is therefore a norm. ■

Proposition 4. The approximate energy inner product $a_h(\cdot, \cdot)$ is uniformly S^h -elliptic.

Proof.

$$\begin{aligned}
a_h(v^h, v^h) &= \sum_{E \in \mathcal{T}^h} \int \int_E \sigma(\Delta v^h)^2 + (1 - \sigma)[(v_{xx}^h)^2 + 2(v_{xy}^h)^2 + (v_{yy}^h)^2] dx dy + \beta \sum_{(x_i, y_i) \in C} v^h(x_i, y_i)^2 \\
&= (1 - \sigma) \left(\sum_{E \in \mathcal{T}^h} \int \int_E (v_{xx}^h)^2 + (v_{xy}^h)^2 + (v_{yy}^h)^2 dx dy + \sum_{(x_i, y_i) \in C} v^h(x_i, y_i)^2 \right) \\
&\quad + \sum_{E \in \mathcal{T}^h} \int \int_E \sigma(\Delta v^h)^2 + (1 - \sigma)(v_{xy}^h)^2 dx dy + (\beta + \sigma - 1) \sum_{(x_i, y_i) \in C} v^h(x_i, y_i)^2 \\
&\geq (1 - \sigma) \|v^h\|_h^2, \quad \text{for } 0 \leq \sigma < 1, \quad \beta \geq 1 - \sigma.
\end{aligned}$$

Since $1 - \sigma$ is positive for $0 \leq \sigma < 1$, $a_h(\cdot, \cdot)$ is uniformly S^h -elliptic. ■

Therefore, the approximate variational principle has a unique solution $u^h \in S^h$. Moreover, because the ellipticity is uniform, Strang's lemma (Theorem A.5) applies, and u^h will converge to the exact solution $u \in V$ as $h \rightarrow 0$, if the approximation is consistent in the sense of equation (A.21). To verify consistency, we apply the patch test, Theorem A.6.

Proposition 5. The square, nonconforming element whose local, quadratic function is defined by (1) and (2) passes the patch test.


Proof. Consider an arbitrary patch of four adjacent elements, all of which share a common node $v_{i,j}$ internal to the patch, as shown in Figure 5. Now, suppose that we impose a constant strain condition on the patch; that is, suppose that we constrain the displacements at all remaining nodes around the periphery of the patch by assigning to them values consistent with the function $\pi_2 \in \Pi^2$, an arbitrary second-degree polynomial. Next, we solve the approximate variational principle (A.17) over the patch domain. This reduces to solving for the unknown displacement at the common unconstrained node $v_{i,j}$ such that it minimizes a quadratic equation. It is a matter of routine algebraic manipulation to show that the displacement obtained will also be consistent with π_2 (we omit the details). In fact, one can show that this is true for the internal, unconstrained nodes of an arbitrary patch of any number of elements whose boundary nodes are made consistent with π_2 . Therefore, the element passes the patch test. ■

Having proved the above propositions, we can now be secure in the fact that our approximate variational principle will provide unique discrete solutions which will converge to the exact solution of the continuous problem as the discretization is made increasingly fine. A limit to the order of convergence that we can expect from our approximation is given by (A.15) — since our element is complete only through quadratics ($k = 2$), we are limited to a convergence rate of order h^2 in displacement ($s = 0$). For a more precise statement, we should take into account the consistency error term in equation (A.20). Nevertheless, we will bypass this complicated analysis because the consistency error is not expected to be large for smooth u which is normally the case when approximating smooth surfaces.

4.4. The Discrete Problem

We are finally ready to derive an explicit form of the discrete problem associated with our

Figure 5. Applying the patch test to four adjacent elements.



approximate variational principle. There are essentially two ways of proceeding. One possibility is to find the Ritz basis functions ϕ_i which are associated with our finite element and which span the space S^h . The basis functions are nonconforming piecewise quadratics with local support, and a basis function associated with each node of the triangulation. We can then use the variational equation (A.18) directly, and write the discrete problem as the linear system of equations analogous to equation (A.13) by computing the matrix coefficients $a_h(\phi_i, \phi_j)$. Unfortunately, the piecewise continuous nature of the basis functions makes them tedious to manipulate, especially near the boundary. We will adopt an alternate approach which altogether avoids the derivation and manipulation of the basis functions. The approach is to solve the approximate variational principle by minimizing the functional $\mathcal{E}_h(\cdot)$ of equation (A.17). Before doing so, however, we make two additional simplifications.

The first simplification involves taking a conservative stance once more. There is no reason to believe that the human visual system is biased in the depth values it assigns such as, for example, making all of them too small or too large. That is to say, we have no reason to assume that there is an external influence on the surface other than that provided by the constraints C , and we should therefore nullify the externally-applied surface force: $g(x, y) \equiv 0$. The linear form (3.4) for the free plate then reduces to

$$f(v^h) = \beta \sum_{(x_i, y_i) \in C} \left(c_{(x_i, y_i)} v^h(x_i, y_i) - \frac{1}{2} c_{(x_i, y_i)}^2 \right). \quad (4)$$

The second simplification involves the choice of a numerical value for the constant σ in our approximate energy inner product $a_h(\cdot, \cdot)$ given by (3). According to the proof of Proposition 4, we are at liberty to choose any value in the range $0 \leq \sigma < 1$, therefore, the simplifying choice $\sigma = 0$ will be made.⁴ Setting $\sigma = 0$ in (3), the energy inner product becomes

$$a_h(u^h, v^h) = \sum_{E \in \mathcal{T}^h} \int_E u_{xx}^h v_{xx}^h + 2u_{xy}^h v_{xy}^h + u_{yy}^h v_{yy}^h dx dy + \beta \sum_{(x_i, y_i) \in C} u^h(x_i, y_i) v^h(x_i, y_i). \quad (5)$$

Thus, according to equation (A.17), we obtain the simplified energy functional

$$\begin{aligned} \mathcal{E}_h(v^h) &= \frac{1}{2} a_h(v^h, v^h) - f(v^h) \\ &= \frac{1}{2} \sum_{E \in \mathcal{T}^h} \int_E (v_{xx}^h)^2 + 2(v_{xy}^h)^2 + (v_{yy}^h)^2 dx dy + \frac{\beta}{2} \sum_{(x_i, y_i) \in C} [v^h(x_i, y_i) - c(x_i, y_i)]^2. \end{aligned} \quad (6)$$

The expression inside the integral will be recognized as the “quadratic variation” expression used by Grimson [1981a].

Since the triangulation over the rectangular region $\bar{\Omega}$ is that of a uniform square grid, it is convenient to impose on the nodes the natural lexicographic indexing scheme implied by Figure 4. We index the nodes by (i, j) for $i = 1, \dots, N_x$ and $j = 1, \dots, N_y$ where N_x and N_y are the number of nodes along the x and y axis respectively. The total number of nodes is $N = N_x \times N_y$. The displacement at node (i, j) is denoted by the variable $v_{i,j}^h$, and all the displacements together are denoted by the vector $\mathbf{v}^h \in \mathbb{R}^N$.

The next step is to express the functional in terms of the node displacements with the help of our element. Inside each element domain E , v^h is a quadratic polynomial given by (1) and (2). Therefore, the second partial derivatives of p^E are constant within E , and are given by

$$\begin{aligned} v_{xx}^h|_E = p_{xx}^E &= 2a = \frac{1}{h^2} (v_{i+1,j}^h - 2v_{i,j}^h + v_{i-1,j}^h); \\ v_{yy}^h|_E = p_{yy}^E &= 2b = \frac{1}{h^2} (v_{i,j+1}^h - 2v_{i,j}^h + v_{i,j-1}^h); \\ v_{xy}^h|_E = p_{xy}^E &= c = \frac{1}{h^2} (v_{i+1,j+1}^h - v_{i,j+1}^h - v_{i+1,j}^h + v_{i,j}^h); \end{aligned}$$

where it is assumed that i, j is the index of the lower left hand node of E . The form of these second derivatives will be recognized as being simply the *finite difference* approximations of order h^2 for the respective derivatives on a uniform, square mesh (see, e.g., [Abramowitz and Stegun, 1965, pg. 884]). Of course, this result is a consequence of our particular choice of finite element, and it will lead to a particularly simple discrete problem. With other elements one cannot expect to obtain finite difference expressions, even for uniform triangulations. Substituting the expressions

⁴Recall that σ is the Poisson constant of the elastic material, so our choice implies that the material does not change in width as it stretches lengthwise. Although this value is not meaningful physically, it is perfectly acceptable mathematically. Aside from a question of convenience, there is further evidence that supports this choice in terms of the optical laws of image formation (see [Grimson, 1982b]).

for the derivatives into $a_h(v^h, v^h)$ in (6) and noting that the area of each element is h^2 , we obtain⁵

$$\begin{aligned}\mathcal{E}_h(\mathbf{v}^h) &= \frac{1}{2} \sum_{E \in \mathcal{T}^h} \int_E (p_{xx}^E)^2 + 2(p_{xy}^E)^2 + (p_{yy}^E)^2 dx dy + \frac{\beta}{2} \sum_{(i,j) \in C} (v_{i,j}^h - c_{i,j}^h)^2 \\ &= \frac{h^2}{2} \sum_{E \in \mathcal{T}^h} [(2a)^2 + 2c^2 + (2b)^2] + \frac{\beta}{2} \sum_{(i,j) \in C} (v_{i,j}^h - c_{i,j}^h)^2 \\ &= \frac{1}{2h^2} \sum_{E \in \mathcal{T}^h} \left[(v_{i+1,j}^h - 2v_{i,j}^h + v_{i-1,j}^h)^2 + 2(v_{i+1,j+1}^h - v_{i,j+1}^h - v_{i+1,j}^h + v_{i,j}^h)^2 \right. \\ &\quad \left. + (v_{i,j+1}^h - 2v_{i,j}^h + v_{i,j-1}^h)^2 \right] + \frac{\beta}{2} \sum_{(i,j) \in C} (v_{i,j}^h - c_{i,j}^h)^2.\end{aligned}$$

We can write the above expression for the functional (aside from the additive constant term) in matrix form as

$$\mathcal{E}_h(\mathbf{v}^h) = \frac{1}{2} (\mathbf{v}^h, \mathbf{A}^h \mathbf{v}^h) - (\mathbf{f}^h, \mathbf{v}^h), \quad (7)$$

where $(\cdot, \cdot): \mathbb{R}^N \times \mathbb{R}^N \mapsto \mathbb{R}$ denotes the familiar Euclidean inner product, and $\mathbf{A}^h \in \mathbb{R}^{NN}$ is a matrix of coefficients. Clearly, equation (7) is the discrete equivalent of the functional (6). For the linear term, we have $\mathbf{f}^h = \beta \mathbf{c}^h$ where $\mathbf{c}^h \in \mathbb{R}^N$ is a vector whose entries associated with constrained displacements are the constraint values $c_{i,j}^h$ and the rest are zero. On the other hand, the matrix $\mathbf{A}^h \in \mathbb{R}^{NN}$ which forms the quadratic term, is a matrix of coefficients which can be broken down as the sum of two matrices: $\mathbf{A}^h = \mathbf{A}_\phi^h + \mathbf{B}^h$. The matrix \mathbf{B}^h is a diagonal matrix whose diagonal entries associated with constrained displacements are equal to β , and the remainder are zero. As is clear from equation (A.13), the entries of the other component \mathbf{A}_ϕ^h can be interpreted as inner products between pairs of basis functions of the finite element space S^h . Since the basis functions have local support, most of these inner products will be zero thereby making \mathbf{A}^h *sparse and banded*. Moreover, since by Propositions 1 and 4, the energy inner product is symmetric and S^h -elliptic, \mathbf{A}^h is a *positive definite, symmetric* matrix. These are important properties from a computational point of view.

To obtain the minimum of $\mathcal{E}_h(\mathbf{v}^h)$ we set to zero its partial derivatives with respect to each of the displacements $v_{i,j}^h$. The minimizing vector of displacements \mathbf{u}^h satisfies the condition

$$\nabla \mathcal{E}_h(\mathbf{u}^h) = \mathbf{A}^h \mathbf{u}^h - \mathbf{f}^h = \mathbf{0},$$


(where ∇ is the gradient operator) which yields a discrete problem in the form of a system of linear equations:

$$\mathbf{A}^h \mathbf{u}^h = \mathbf{f}^h, \quad (8)$$

where the entries of \mathbf{A}^h are given by

⁵We also assume for simplicity that all constraints $c_{(x_i, y_i)}$ coincide with nodes (i, j) in \mathcal{T}^h . Hence, we will denote the constraints by $c_{i,j}^h$.

Figure 6. Computational molecules associated with the discrete problem.



$$A^h = \left[\frac{\partial^2 \mathcal{E}_h(\mathbf{u}^h)}{\partial u_{i,j} \partial u_{k,l}} \right], \quad 1 \leq i, k \leq N_x, \quad 1 \leq j, l \leq N_y.$$

From this expression, A^h will be recognized as the *Hessian matrix* of the functional \mathcal{E}_h (see, e.g., [Luenberger, 1973]).

Although the evaluation of the Hessian matrix entries is routine for interior nodes, it is tedious due to the special cases for the elements around the boundaries of the region $\bar{\Omega}$. We

omit the details and give the final result in terms of a set of *computational molecules* which are illustrated in Figure 6 in relation to the lower left hand edge of $\bar{\Omega}$ whose boundary is indicated by bold links. Obviously, computational molecules for the remaining edges are appropriate rotations of those shown. The particular computational molecule associated with a node specifies the (nonzero) coefficients of the equation for that node. For example, the equation for the displacement at a node (i, j) in the interior of the region (indicated by the double circle in the topmost computational molecule in Figure 6) is

$$\begin{aligned} & -\frac{8}{h^2}(v_{i-1,j}^h + v_{i+1,j}^h + v_{i,j-1}^h + v_{i,j+1}^h) \\ & + \frac{2}{h^2}(v_{i-1,j-1}^h + v_{i+1,j-1}^h + v_{i-1,j+1}^h + v_{i+1,j+1}^h) \\ & + \frac{1}{h^2}(v_{i-2,j}^h + v_{i+2,j}^h + v_{i,j-2}^h + v_{i,j+2}^h) \\ & + \frac{20}{h^2}v_{i,j}^h + \beta v_{i,j}^h = \beta c_{i,j}. \end{aligned} \tag{9}$$

The terms involving β are present only if there is a constraint $c_{i,j}$ at that node.

The sparseness of Λ^h is evident from the above equation — matrix rows associated with interior nodes have only 13 nonzero entries, while rows associated with nodes near the boundaries have even fewer. Also, note that the computational molecule for the center of the region is a factor of h^2 (due to the elemental area) times the finite difference approximation of order h^2 for the biharmonic operator [Abramowitz and Stegun, 1965, pg. 885] that is associated with the Euler-Lagrange equation for our variational principle. This is an expected consequence of our particular choice of element which yielded finite difference approximations for the second partial derivatives of v^h . Moreover, aside from multiplicative constants, the same molecules were obtained by Grimson [1981a] in the specification of a (conjugate gradient) mathematical programming algorithm. As was previously argued however, the finite element method is richer in that it systematically suggests many alternative, less-restrictive triangulations, as well as more general local representations for surfaces.

5. MULTI-LEVEL SURFACE RECONSTRUCTION

As we have seen, the application of the finite element method to a well-posed quadratic variational principle, such as the one on which our surface approximation problem is based, inevitably leads to an equivalent discrete problem which takes the form of a linear system of algebraic equations. The matrix of coefficients of this nonsingular system is symmetric, positive definite, sparse, and banded. Computing the most consistent approximating surface now amounts to solving this system and, in this chapter, we adopt an efficient hierarchical algorithm to perform this task. We will proceed to develop the algorithm and to demonstrate its performance it on constraints from stereopsis.

5.1. Possible Methodologies for Solving the Discrete Problem

The solution of linear systems is a very important problem in numerical analysis and the many techniques which have been developed fall into essentially two broad classes — direct methods which yield the solution in a finite number of steps, and iterative methods which typically converge asymptotically to the solution (see, e.g. [Dahlquist and Björck, 1974] or [Gladwell and Wait, 1979]).

Direct methods include matrix inversion methods such as *Gaussian elimination* and *LU decomposition*. Although widely used for solving finite element equations, they usually do not exploit the sparseness and bandedness of the system matrix because, during the inversion process, the sparse matrix is transformed into a full one.⁶ Consequently, all the elements of the matrix must be stored. Moreover, direct methods are typically global and sequential algorithms, which makes them unsuitable models for neurally-based visual processes.

On the other hand, the class of *iterative methods* readily gives rise to biologically-feasible algorithms. Examples in this class are *relaxation methods* such as *Jacobi relaxation*, *Gauss-Seidel relaxation* and *successive overrelaxation*, as well as *gradient methods* such as *gradient descent* and the *conjugate gradient method*. Iterative methods exploit the sparseness of the matrix inasmuch as they do not modify its elements from one iteration to the next. Therefore, only the relatively few nonzero matrix elements need be stored. Owing to the sparseness and banded structure of the matrix, iterative methods require local-support computations only, and in certain forms such as Jacobi relaxation and gradient methods the computations can be performed in parallel. Because iterative methods in general and relaxation in particular are fundamental to the ensuing discussion, an introduction to some of the relevant mathematics of this class of techniques is included in Appendix B for convenience.

The algorithms we are contemplating are to be executed by computational mechanisms in the form of networks of many simple processors, such as neurons, which are directly connected only to near neighbors. Due to the myopic nature of the processors, global interactions can take place only indirectly, through the iterative process, by an incremental propagation of information. Normally, the network is large and since this is reflected in the size of the linear system, we anticipate that a vast number of iterations will be required for any relaxation or gradient method to converge. Typically, the number of iterations will be on the order of N^m , where N is the dimension of the matrix, and m is the highest order of partial derivatives present in the energy inner product, which in our case is two. Grimson's [1981a] formulation of surface interpolation as a problem in mathematical programming naturally led him to the choice of a gradient method for its solution and, not unexpectedly, disappointingly slow convergence rates were observed due to the large size of the images typically encountered.

⁶For a positive definite symmetric matrix, the inverse matrix remains banded, but is no longer sparse within the band. The inverse matrix is the discrete Green's function for our problem, which in general has global support over $\bar{\Omega}$.

Recently, a class of iterative techniques called *multi-grid methods* have seen increased application to the numerical solution of boundary value problems for which they achieve convergence in essentially order N number of operations [Brandt, 1977a, 1977b; Brandt and Dinar, 1979]. This spectacular improvement results from the use of a hierarchy of grids to increase the efficiency of the global propagation of information. Multi-grid algorithms feature both intra-grid and inter-grid processes. Typically, the intra-grid processes are relaxation iterations, while the inter-grid processes are local polynomial interpolations. Therefore the multi-grid algorithms are, in principle, biologically feasible. A final issue which speaks in favor of adopting them to vision is the intrinsic multi-level structure of the earliest stages of the visual system itself and, as we argued in the introduction, the apparent need to maintain this structure at least to the level of the $2\frac{1}{2}$ -D sketch.

We therefore advocate a hierarchical approach to surface reconstruction, which we will develop initially in the context of the Marr-Poggio stereo theory whose clear multi-level structure provides ample motivation. At the heart of the proposed scheme lies a multi-grid algorithm adapted to the fast solution of a hierarchy of discrete thin plate surface approximation problems. In the following sections, we present the underlying theory and build up a detailed description of the multi-level algorithm.

5.2. The Multi-Level Equations

As we have stated, the stereo module generates sparse depth information over a range of resolutions. The information at any particular scale can be thought of as a set of constraints which, at that level, define a well-posed, discrete surface approximation problem. It is natural then to view our surface reconstruction problem as the solution of a hierarchy of such discrete problems. The discretizations are performed in the usual way by introducing a sequence of finite element spaces S^{h_1}, \dots, S^{h_L} over the rectangular domain $\bar{\Omega}$, where L is the number of levels and $h_1 > \dots > h_L$ are the fundamental lengths of the elements at each level. In the familiar notation, we will denote the functions which are members of the finite element spaces by (*italic face*) $v^{h_k} \in S^{h_k}$, and the parameters (i.e., the nodal displacements) which define these functions according to (A.11) by (**bold face**) vectors $\mathbf{v}^{h_k} \in \mathbb{R}^{N^{h_k}}$, where N^{h_k} is the dimension of S^{h_k} . The hierarchy of problems is then given by the sequence of L linear systems (see equation (4.8)) of the form

$$\mathbf{A}^{h_k} \mathbf{u}^{h_k} = \mathbf{f}^{h_k}, \quad 1 \leq k \leq L, \quad (1)$$

whose discrete solutions $\mathbf{u}^{h_k} \in \mathbb{R}^{N^{h_k}}$ for $1 \leq k \leq L$ define a sequence of functions $u^{h_k} \in S^{h_k}$ which constitute the hierarchy of full surface representations.

Although, in theory, there need be no restriction in the relationship of element sizes from one level to the next, a number of practical implementation-related issues point towards the subdivision of each square element domain on a given level S^{h_k} into four identical element domains

on the next finer level $S^{h_{k+1}}$; that is, we choose $h_k = 2h_{k+1}$. Consequently, S^{h_k} will be a subspace of $S^{h_{k+1}}$, and the implementation of the inter-level processes is simplified substantially. Moreover, the 2:1 ratio is a natural one in view of the spatial-frequency bandpass channels in early vision whose center frequencies are spaced approximately one octave apart, the spatial resolution of a channel being about twice that of the immediately coarser one [Wilson and Giese, 1979]. Finally, the choice can be shown to be near optimal in terms of the *multi-grid convergence rate* [Brandt, 1977a, pg. 353]. Since the triangulation of $\bar{\Omega}$ associated with our simple elements is a uniform grid of square element domains, the 2:1 ratio implies that in scanning along the x or y directions, every second node of a grid coincides with a node of the next coarser grid and, furthermore, that the number of nodes is related from one level to the next by $N^{h_{k-1}} = \frac{1}{4}N^{h_k}$. Therefore, the total amount of space required to maintain all of the representations is bounded by $N^{h_L}(1 + \frac{1}{4} + \frac{1}{16} + \dots) = \frac{4}{3}N^{h_L}$; i.e., it is only a small fraction more than that required for the finest grid.

One can think of several possibilities for exploiting the hierarchy of discrete problems to increase the convergence rate of the iterative process. Perhaps the first idea that comes to mind is to solve the system at the coarsest level, which can be done very quickly, and use the solution as an initial approximation in the iterative solution of the next finer level, proceeding in this manner to the finest level. This is an effective acceleration strategy which is almost as old as the idea of relaxation itself [Southwell, 1946]. Although it is suitable for obtaining a single accurate solution at the finest level, it cannot generate solutions having the finest-level accuracy over the hierarchy of coarser levels, since the approximation error increases as the elements become larger. This is undesirable from the point of view of our surface reconstruction problem. Here we require that the accuracy of the finest-resolution surface be maintained throughout the coarser surface descriptions. This will guarantee that the shape of the surface will be consistent over the hierarchy of representations.

The stipulation that accuracy be maintained is further motivated by psychophysical studies into the phenomenon of *visual hyperacuity* (see, e.g., [Westheimer, 1977; Westheimer and McKee, 1975, 1977]). Related computational studies indicate that, in principle, sharp, well-defined intensity edges can be localized to high (sub-receptor separation) accuracies from the $\nabla^2 G$ convolution values through a process of *spatiotemporal interpolation* [Barlow, 1979; Marr *et al.*, 1980; Hildreth, 1980; Crick *et al.*, 1981; Fahle and Poggio, 1981]. Consequently, it seems that although the depth constraints arising from the larger channels in stereopsis represent coarser spatial samplings of the scene, excluding erroneous matches, the samples may provide highly accurate range information.

The only way that consistent accuracy can be maintained throughout the hierarchy of full surface representations is to allow the coarser levels access to the high-resolution information in the finer levels. The multi-grid algorithm provides such a flow. The hierarchy of levels cooperate, through a bi-directional flow of information, to simultaneously generate multiple, equally-accurate

surface representations, and do so with much less computational effort than would be expended in solving the finest-level system in isolation. To see how this is accomplished, we will initially consider only two levels, a fine one and a coarse one, associated with the finite element spaces S^h and S^{2h} respectively. Suppose that by some iterative process we obtain an approximate solution \mathbf{v}^{2h} to the coarse level system $\Lambda^{2h}\mathbf{u}^{2h} = \mathbf{f}^{2h}$, which is then interpolated⁷ to the fine level where it becomes the initial approximation \mathbf{v}^h :

$$\mathbf{v}^h \leftarrow \mathbf{I}_{2h \Rightarrow h} \mathbf{v}^{2h}. \quad (2)$$

The mapping $\mathbf{I}_{2h \Rightarrow h}: S^{2h} \mapsto S^h$ denotes interpolation from the coarse space to the fine space. Normally, \mathbf{v}^h will require substantial improvement.

Let \mathbf{u}^h be the solution to the fine level system $\Lambda^h\mathbf{u}^h = \mathbf{f}^h$. Then we can define the error vector in a given approximation \mathbf{v}^h by $\mathbf{e}^h = \mathbf{u}^h - \mathbf{v}^h$. Clearly, if \mathbf{e}^h could be computed, it could be added as a correction to \mathbf{v}^h , thereby giving us the desired solution. But because the computation of \mathbf{e}^h would take about as much effort as computing \mathbf{u}^h itself, doing so would not be helpful. On the other hand, if we could somehow approximate the error function e^h by a function e^{2h} in the coarse space S^{2h} , such an approximation can be obtained quickly due to the fact that the coarse space has only one quarter the dimensionality of the fine space. Such an approximation is generally not possible, however, because e^h , having been generated by an interpolation from the coarse grid solution, is certain to have large fluctuations with wavelengths less than $4h$. These *high-frequency* Fourier components could not be approximated on the coarse grid because there they would alias as lower-frequency components. Before a meaningful approximation to the error can be obtained on the coarse grid, the high-frequency components must be eliminated; that is to say, the error function e^h must be *smoothed*.

Since smoothing is inherently a local operation, it should not be surprising that local iterative methods, inefficient as they are in obtaining solutions, are very efficient at smoothing the error function. In particular, although relaxation generally requires very many iterations to eliminate the global, low-frequency components of the error, it only takes a few iterations to eliminate the local, high-frequency components. This behavior can be predicted mathematically by a *local Fourier analysis* of the given iterative method [Brandt, 1977a]. The analysis involves a local Fourier expansion of the error function followed by an examination of the effect that a single iteration has on the amplitudes of each component. An important quantity which is obtained through this analysis is the *smoothing factor* $\bar{\mu}$ of the iterative scheme, which is defined as the worst (i.e., the largest) amplification of a high-frequency component of the error from one iteration to the next. As an example, in Appendix C we carry out a local Fourier analysis of the appropriate Gauss-Seidel scheme for our discrete problem, and show that $\bar{\mu} \approx 0.8$. This implies that, for our problem, ten Gauss-Seidel iterations on the fine grid are sufficient to reduce the high-frequency

⁷Lagrange interpolation of a suitable order may be used.

components of e^h by approximately an order of magnitude. A more effective *weighted* Jacobi relaxation scheme, which is also suitable for our problem and gives a $\bar{\mu} = 0.549$, is described in [Brandt, 1977a, pg. 342].⁸

Once the error has been smoothed, it may be inexpensively approximated on the coarse grid. The equation for e^h on the fine grid is the *residual equation*

$$\Lambda^h e^h = r^h, \quad \text{where} \quad r^h = f^h - \Lambda^h v^h \quad (3)$$

is called the *residual* of the approximation v^h . The approximation to this equation on the coarser grid is

$$\Lambda^{2h} e^{2h} = \mathcal{I}_{h \Rightarrow 2h} r^h,$$

where the mapping $\mathcal{I}_{h \Rightarrow 2h}: S^h \mapsto S^{2h}$ is an "interpolation" from the fine space to the coarse space. Because $S^{2h} \subset S^h$, the mapping can be a simple *injection* or some local averaging of node displacements from the fine grid to the coarse. After e^{2h} is computed, a better approximation to v^h may be obtained by interpolating the coarse grid correction back to the fine grid; that is, by making the replacement:

$$v^h \leftarrow v^h + \mathcal{I}_{2h \Rightarrow h} e^{2h}.$$

This correction practically annihilates the smooth part of the error e^h .

Brandt [1977a, 1977b] calls the foregoing scheme a *correction scheme* in view of the fact that the function computed on the coarse grid is the error function; that is, the correction to the fine grid approximation. The correction scheme is easy to implement, but it is unsuitable for our surface reconstruction problem because instead of an error function e^{2h} , we require that the function computed in the coarse space be a function u^{2h} which represents explicitly the distances to surfaces in the scene. This may be accomplished by a reformulation of the correction scheme equations which converts them into those of the related *full approximation scheme*.

First, we rewrite (3) in the equivalent form

$$\Lambda^h (v^h + e^h) - \Lambda^h v^h = r^h,$$

which may be approximated by the corresponding coarse-grid equation

$$\Lambda^{2h} \left(\mathcal{I}_{h \Rightarrow 2h} v^h + e^{2h} \right) - \Lambda^{2h} \left(\mathcal{I}_{h \Rightarrow 2h} v^h \right) = \mathcal{I}_{h \Rightarrow 2h} r^h.$$

Defining a new function u^{2h} in S^{2h} by the nodal displacement vector $u^{2h} = \mathcal{I}_{h \Rightarrow 2h} v^h + e^{2h}$, we obtain the coarse level system

⁸Brandt proposes this scheme for solving biharmonic boundary value problems. The scheme is also appropriate for our surface approximation problem which is in essence a biharmonic problem in view of the associated Euler-Lagrange equation.

$$\mathbf{A}^{2h} \mathbf{u}^{2h} = \mathbf{g}^{2h}, \quad \text{where} \quad \mathbf{g}^{2h} = \mathbf{A}^{2h} \left(\mathbf{I}_{h \Rightarrow 2h} \mathbf{v}^h \right) + \mathbf{I}_{h \Rightarrow 2h} \mathbf{r}^h. \quad (4)$$

It is natural to interpret (4) as the original coarse-level system $\mathbf{A}^{2h} \mathbf{u}^{2h} = \mathbf{f}^{2h}$ but with a right-hand side which has been modified using information from the fine grid so as to maintain the fine-grid accuracy in the coarse-grid function \mathbf{u}^{2h} . Thus, \mathbf{g}^{2h} is an *estimate of the local truncation error* on the coarse level relative to the fine level (see [Brandt, 1977b, pg. 284]).

Once the solution \mathbf{u}^{2h} of equation (4) is available, we can write $\mathbf{e}^{2h} = \mathbf{u}^{2h} - \mathbf{I}_{h \Rightarrow 2h} \mathbf{v}^h$ as the desired coarse-grid approximation to the fine-grid error, and the approximation on the fine level can be corrected by the replacement

$$\mathbf{v}^h \leftarrow \mathbf{v}^h + \mathbf{I}_{2h \Rightarrow h} (\mathbf{u}^{2h} - \mathbf{I}_{h \Rightarrow 2h} \mathbf{v}^h). \quad (5)$$

Note that since $\mathbf{I}_{2h \Rightarrow h} \mathbf{I}_{h \Rightarrow 2h} \mathbf{v}^h \neq \mathbf{v}^h$ the replacements given by (2) and (5) are *not* equivalent. Since \mathbf{u}^{2h} is a low-frequency correction, the replacement indicated by (2) would destroy the high-frequency components of \mathbf{v}^h whereas the replacement indicated by (5) preserves them.

How do we solve the coarse-grid equation (4)? The obvious answer is: by relaxation iterations on the coarse grid, and with the help of corrections obtained from still coarser grids. Thus, in a straightforward recursive fashion, we can extend the above two-level equations to any number of levels. In view of (4) and the fact that the residual for the level k equations is given by

$$\mathbf{r}^{hk} = \mathbf{g}^{hk} - \mathbf{A}^{hk} \mathbf{u}^{hk}, \quad (6)$$

the *multi-level equations* for L levels are given by

$$\mathbf{A}^{hk} \mathbf{u}^{hk} = \mathbf{g}^{hk} \quad \text{for } 1 \leq k \leq L, \quad (7)$$

where

$$\begin{aligned} \mathbf{g}^{h_L} &= \mathbf{f}^{h_L}; \quad \text{and} \\ \mathbf{g}^{hk} &= \mathbf{A}^{hk} \left(\mathbf{I}_{h_{k+1} \Rightarrow h_k} \mathbf{u}^{h_{k+1}} \right) + \mathbf{I}_{h_{k+1} \Rightarrow h_k} (\mathbf{g}^{h_{k+1}} - \mathbf{A}^{h_{k+1}} \mathbf{u}^{h_{k+1}}), \quad \text{for } 0 \leq k \leq L-1. \end{aligned} \quad (8)$$

Note that the original, right-hand side \mathbf{f}^{h_k} of the k^{th} level problem occurs only on the finest level L . The right-hand sides of the coarser levels have been modified in order that the finest level accuracy be properly maintained throughout; that is to say, in order for the solutions \mathbf{u}^{h_k} to coincide: $\mathbf{u}^{h_1} = \mathbf{I}_{h_2 \Rightarrow h_1} \mathbf{u}^{h_2} = \dots = \mathbf{I}_{h_2 \Rightarrow h_1} \dots \mathbf{I}_{h_L \Rightarrow h_{L-1}} \mathbf{u}^{h_L}$. Analogously to the two-grid case, we can interpret the difference of the original and the corrected right-hand sides, $\mathbf{f}^{h_k} - \mathbf{g}^{hk}$, as an estimate of the *local truncation error* of level k relative to the finest level.

5.3. Multi-Level Surface Reconstruction Algorithms

We have motivated the multi-level approach to surface reconstruction and described in a quantitative manner its basic components — the intra-level relaxation processes, and the inter-level interpolation processes. It now remains to show how to bring the components together into an

algorithm for solving the multi-level equations (7) and (8). Several schemes have been proposed [Brandt, 1977a, 1977b; Brandt and Dinar, 1979]. We will describe one which is appropriate in terms of our surface approximation problem.⁹ Before defining the *full multi-level surface reconstruction algorithm*, we will define its main procedure, the *multi-level surface reconstruction cycle*.

The multi-level surface reconstruction cycle starts at the (currently) finest level l , making several cycles to the coarser levels $k = l-1, l-2, \dots, 1$, until a hierarchy of surface representations which are as accurate as is possible in the S^{h_l} space is obtained. Let ϵ_k denote a tolerance for solving the equations on level k . ξ and η are *switching parameters* which are given appropriate values below.

Algorithm 1 — Multi-Level Surface Reconstruction Cycle

Step 1 — initialize the finest level l .

Set the right hand side of the level l problem $A^{h_l} u^{h_l} = g^{h_l}$ to the original right hand side: $g^{h_l} \leftarrow f^{h_l}$. Introduce the initial approximation $v^{h_l} \leftarrow v_0^{h_l}$. Set $\epsilon_l \leftarrow 0$,¹⁰ and $k \leftarrow l$.

Step 2 — start a new operation level k .

Set $e_k^{\text{old}} \leftarrow \infty$.

Step 3 — perform a relaxation iteration.

Perform a relaxation iteration for the equation $A^{h_k} u^{h_k} = g^{h_k}$ and concurrently compute some norm of the residual given by (6), $e_k \leftarrow \|r^{h_k}\|$.

Step 4 — test the convergence and its rate.

If $e_k \leq \epsilon_k$, then convergence has been obtained at the current operation level; go to Step 6. If $k = 1$, go to Step 3. If $e_k \leq \eta e_k^{\text{old}}$ then the convergence rate is still satisfactory, set $e_k^{\text{old}} \leftarrow e_k$ and go to Step 3; otherwise the convergence rate is slow so go to Step 5.

Step 5 — transfer to coarser level.

Introduce as the first coarse-level approximation the function $v^{h_{k-1}}$ defined by the nodal displacements

$$v^{h_{k-1}} \leftarrow \bigcup_{h_k \Rightarrow h_{k-1}} v^{h_k}.$$

Set the right-hand side of the coarser level problem $A^{h_{k-1}} u^{h_{k-1}} = g^{h_{k-1}}$ to

$$g^{h_{k-1}} \leftarrow A^{h_{k-1}} v^{h_{k-1}} + \bigcup_{h_k \Rightarrow h_{k-1}} (g^{h_k} - A^{h_k} v^{h_k})$$

⁹Brandt refers to it as the *accommodative, full multi-grid, full approximation scheme* algorithm [Brandt and Dinar, 1979].

¹⁰This value for ϵ_l is temporary. A realistic value is introduced in Step 5.

(in view of equation (8)). Set the tolerance $\epsilon_{k-1} \leftarrow \xi e_k$. Concurrently with the computation of $\mathbf{g}^{h_{k-1}}$, compute the norm of the local truncation error $\mathbf{f}^{h_{k-1}} - \mathbf{g}^{h_{k-1}}$ using the same norm as in Step 3. If $k = l$ set $\epsilon_l \leftarrow \frac{1}{4} \|\mathbf{f}^{h_{k-1}} - \mathbf{g}^{h_{k-1}}\|$.¹¹ Finally, set $k \leftarrow k - 1$ and go to Step 2.

Step 6 — use converged solution \mathbf{u}^{h_k} to correct finer level.

If $k < l$, make the correction (in view of equation (5))

$$\mathbf{v}^{h_{k+1}} \leftarrow \mathbf{v}^{h_{k+1}} + \mathbf{I}_{h_k \Rightarrow h_{k+1}} (\mathbf{u}^{h_k} - \mathbf{I}_{h_{k+1} \Rightarrow h_k} \mathbf{v}^{h_{k+1}});$$

set $k \leftarrow k + 1$ and go to Step 2. Otherwise, if $k = l$ end.

The relaxation operation in Step 3 can, in principle, be based on any one of the iterative methods described in Appendix B, but is usually a simple Jacobi or Gauss-Seidel iteration. For our surface reconstruction problem, in view of equation (4.9) and the discussion in Appendix B, the Jacobi relaxation iteration in the interior of $\bar{\Omega}$ is given by

$$\begin{aligned} v_{i,j}^{(i+1)} = \frac{1}{\frac{20}{h^2} + \beta} & \left[\frac{8}{h^2} (v_{i-1,j}^{(i)} + v_{i+1,j}^{(i)} + v_{i,j-1}^{(i)} + v_{i,j+1}^{(i)}) \right. \\ & - \frac{2}{h^2} (v_{i-1,j-1}^{(i)} + v_{i+1,j-1}^{(i)} + v_{i-1,j+1}^{(i)} + v_{i+1,j+1}^{(i)}) \\ & \left. - \frac{1}{h^2} (v_{i-2,j}^{(i)} + v_{i+2,j}^{(i)} + v_{i,j-2}^{(i)} + v_{i,j+2}^{(i)}) + \beta g_{i,j} \right], \end{aligned} \quad (9)$$

while the Gauss-Seidel relaxation iteration is given by

$$\begin{aligned} v_{i,j}^{(i+1)} = \frac{1}{\frac{20}{h^2} + \beta} & \left[\frac{8}{h^2} (v_{i-1,j}^{(i+1)} + v_{i+1,j}^{(i)} + v_{i,j-1}^{(i+1)} + v_{i,j+1}^{(i)}) \right. \\ & - \frac{2}{h^2} (v_{i-1,j-1}^{(i+1)} + v_{i+1,j-1}^{(i+1)} + v_{i-1,j+1}^{(i)} + v_{i+1,j+1}^{(i)}) \\ & \left. - \frac{1}{h^2} (v_{i-2,j}^{(i+1)} + v_{i+2,j}^{(i)} + v_{i,j-2}^{(i+1)} + v_{i,j+2}^{(i)}) + \beta g_{i,j} \right], \end{aligned} \quad (10)$$

where we have suppressed the superscripts h_k indicating the level, and have instead introduced the bracketed superscripts which indicate the iteration number. Analogous formulas for the boundary points can be derived from the computational molecules associated with the boundary. The norm computed in Steps 3 and 5 can be the discrete L_2 or L_∞ norm. In the case of Gauss-Seidel relaxation, it is quicker to compute the *dynamic norm*, as the iteration progresses, rather than the static norm (see, e.g. [Brandt, 1977b, pg. 286]).

An important feature of the multi-level algorithm is that the local Fourier analysis, in addition to providing a prediction of the convergence rate, enables one to predict near-optimal values for the switching parameters. It turns out that the actual values assigned to the switching parameters are not critical, and that good values are $\xi = 0.2$ and $\eta = \bar{\mu}$, where $\bar{\mu}$ is the smoothing

¹¹The constant $\frac{1}{4}$ is the value of 2^{-p} (see [Brandt, 1979, pg. 65]), where $p = 2$ is the approximation order of the second partial derivatives of the energy inner product that is achieved by our quadratic element.

factor of the relaxation method used in Step 3 of the algorithm (see [Brandt, 1977b, pg. 290]). The order of the interpolation operators is determined by the problem itself; i.e. the order of derivatives in the energy inner product. For the coarse-to-fine interpolation $I_{h_k \Rightarrow h_{k+1}}$ in Step 6, the natural second-degree interpolation of the element polynomial p^E may be used. On the other hand, simple injections perform well for the fine-to-coarse transfers $I_{h_k \Rightarrow h_{k-1}}$ in Step 5 and $I_{h_{k+1} \Rightarrow h_k}$ in Step 6.

Having defined the multi-level cycle which starts at the finest level l and cycles through the coarser levels, we will employ it as a procedure within a more general, full multi-level surface reconstruction algorithm. We now think of the level l of the cycling algorithm as the *currently finest level*; i.e., the finest level for which an approximate solution has already been computed by the multi-level cycle. The full algorithm works in the opposite direction, the currently finest level progressing from the coarsest level $l = 1$, to the finest level $l = L$.

Algorithm 2 — Full, Multi-Level Surface Reconstruction Algorithm

Step 1 — solve the coarsest-grid equation.

Compute by relaxation an approximate solution u^{h_1} to the coarsest-grid equation $A^{h_1} u^{h_1} = f^{h_1}$. Set $l \leftarrow 1$.

Step 2 — set a new finest level l .

If $l = L$ stop. Otherwise increment the currently finest level $l \leftarrow l + 1$, and set the first approximation on the new level to be the function $v_0^{h_l}$ defined by nodal displacements $v_0^{h_l} = I_{h_{l-1} \Rightarrow h_l} u^{h_{l-1}}$.

Step 3 — perform a multi-level cycle.

Invoke Algorithm 1 and when it ends, go to Step 2.

Note that the solution in Step 1 will be performed quickly because S^{h_1} has relatively few dimensions. In Step 3, each time Algorithm 1 terminates at level l , we have obtained a hierarchy of l representations whose accuracy is the best possible on level l . The currently finest approximation is then interpolated to the next finer level until the finest level L is reached. Brandt recommends a somewhat higher-order interpolation for the initial interpolation $I_{h_{l-1} \Rightarrow h_l}$ in Step 2. Third order Lagrangian interpolation seems adequate for our surface interpolation problem, as we will see from the demonstrations in the next section.

Algorithm 1 is *accommodative* in that it makes internal checks, based on the computation of norms, to determine when to switch levels. For many types of problems, accommodative algorithms behave in a fairly fixed manner, performing a similar number of iterations on each level before switching. It is then possible to avoid computing the dynamic-residual norm in Step 3 of Algorithm 1, and to preassign a fixed flow. A switch to the coarser level $S^{h_{k-1}}$ is made after n_c iterations have been performed at level S^{h_k} . Analogously, a switch to the finer level $S^{h_{k+1}}$ is made after n_f iterations have been performed on level S^{h_k} since the last return from the finer

level. n_c depends on the smoothing factor, a good choice being $n_c = \log .1 / \log \bar{\mu}$. Sometimes n_f varies from level to level. For a more extensive discussion see [Brandt, 1979, pg. 68–69]. *Fixed algorithms* are to be preferred for parallel implementations in general, and from a biological point of view in particular.

In order to evaluate the performance of the multi-level surface reconstruction algorithm, we define a unit of computation called a *work unit* which is the amount of computation required to perform one relaxation iteration on the finest level L . It is roughly equal to wN^{h_L} , where w is the number of operations required to compute the residual at each node¹² and N^{h_L} is the number of nodes at the finest level. Since there are one quarter as many nodes on level $k - 1$ as there are on level k , only $1/4^i$ work unit is required to perform a relaxation iteration on level $L - i$. The proportionate amount of computation done on coarser grids thus diminishes very rapidly. Although, for accommodative algorithms, it is difficult to predict the total number of operations consumed by the inter-level processes, it normally turns out to be considerably less than the total inter-level process computation, and is therefore usually ignored (see [Brandt, 1977a, section 6.2]).

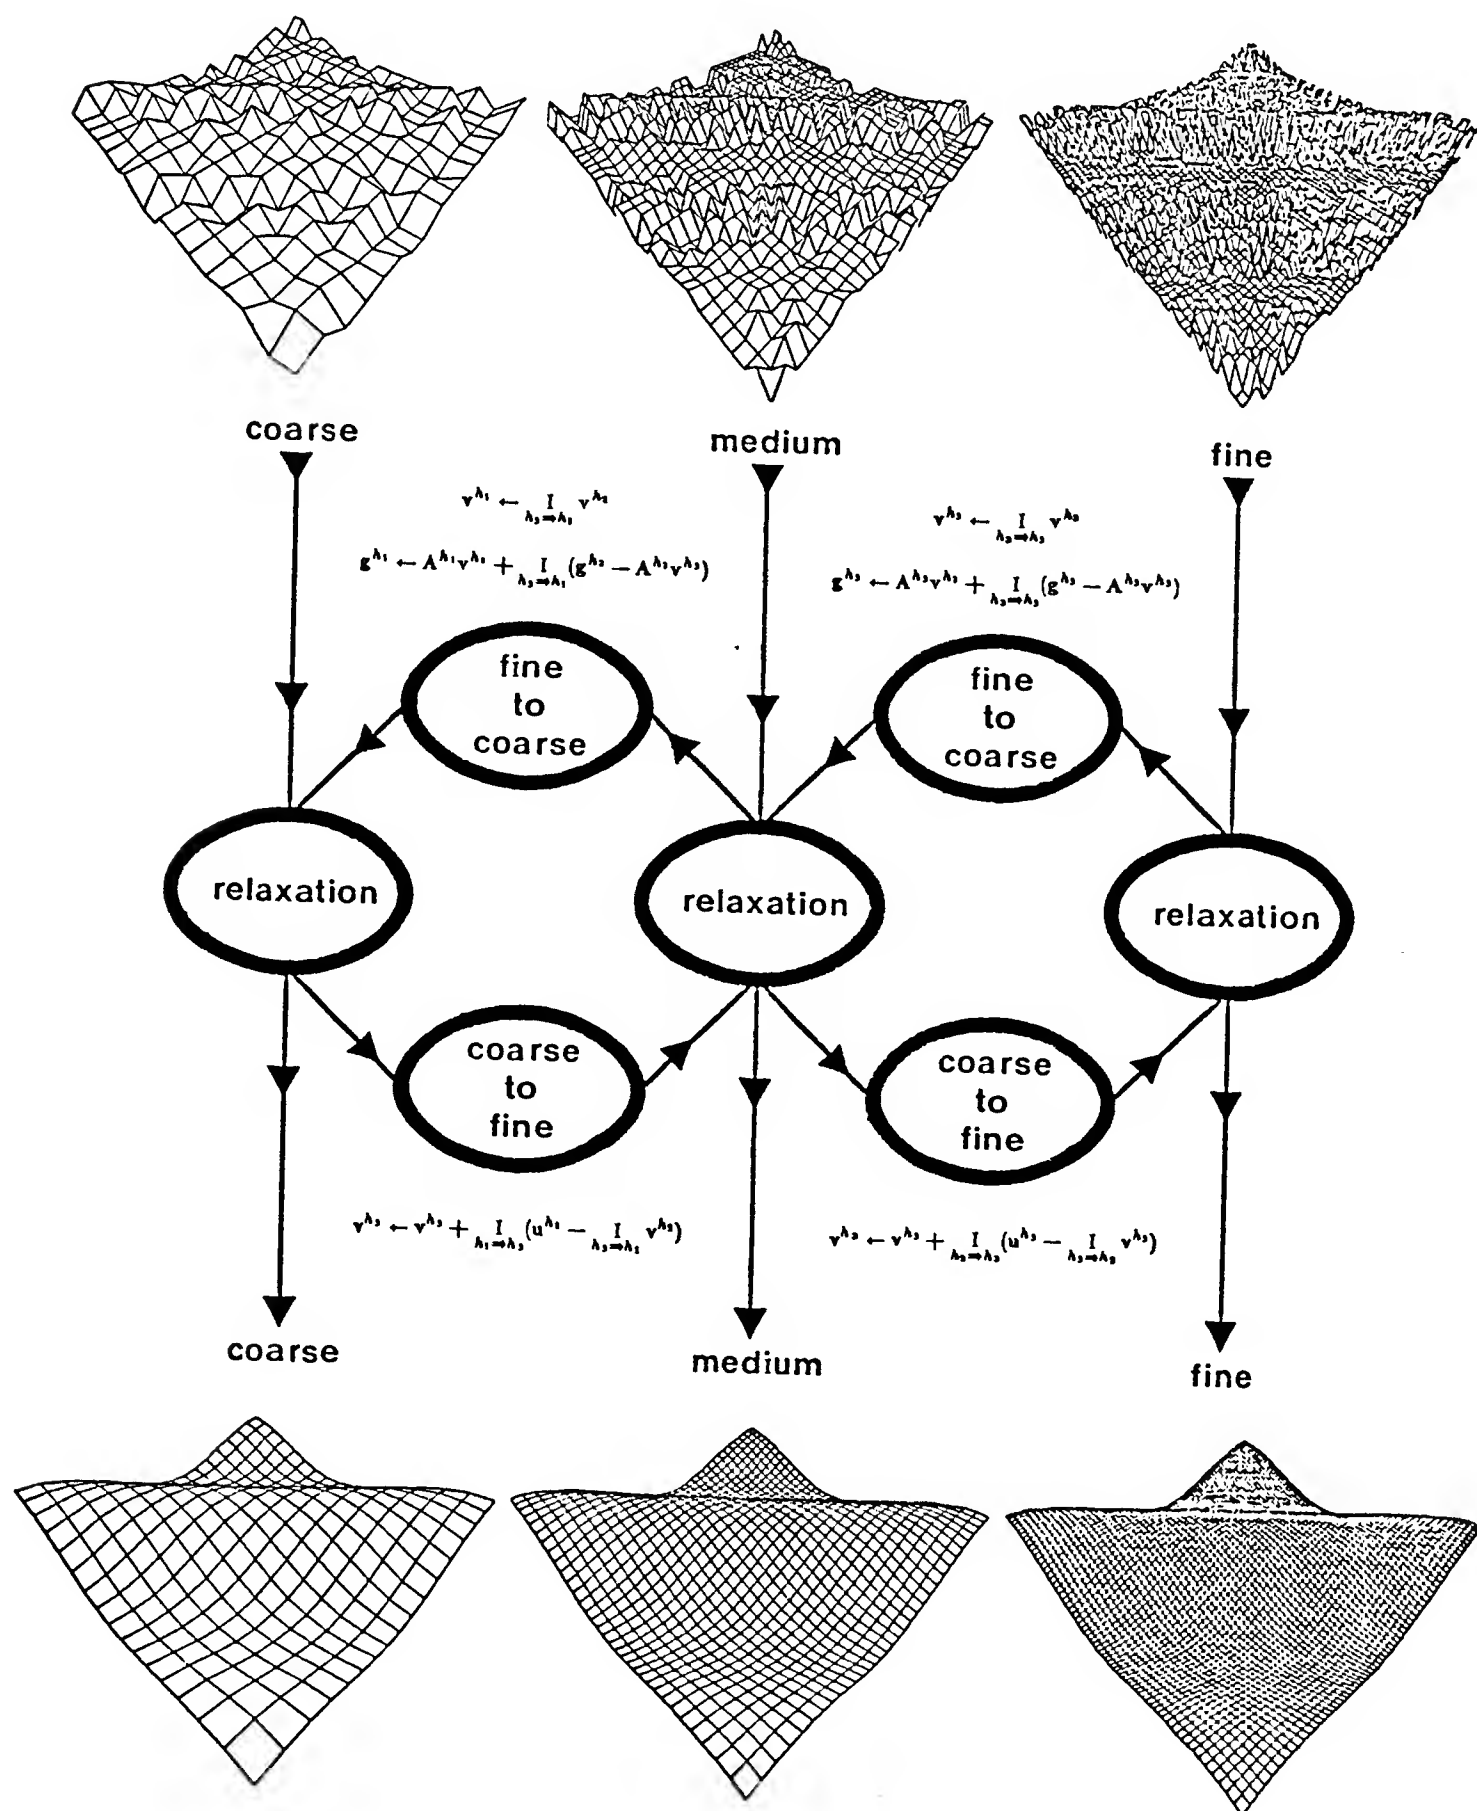
A final issue that we have not yet considered in quantitative terms is the choice of appropriate values for the (spring) constant β . In the mathematical and, in particular, in the finite element literature, the constraint term $\mathcal{E}_5(v)$ of equation (2.2) is known as a *penalty function* (see, e.g., [Courant, 1943], [Babuska, 1973], [Strang and Fix, 1973], [Zienkiewicz, 1974]). The incorporation of penalty functions into variational principles is a standard way of approximately satisfying essential boundary conditions by converting them into appropriate natural boundary conditions which may be handled straightforwardly by the finite element method. Penalty functions are particularly useful when the essential boundary conditions in question are complicated, or when only their approximate satisfaction is desired, as in the case of visual surface approximation. An optimal value for β can be derived through the following considerations. Let w be the solution to our surface approximation problem, which interpolates the known depth points. As usual, $u \in \mathcal{H}^2(\Omega)$ denotes the exact solution to the variational principle (2.2), including the penalty term \mathcal{E}_5 , and $u^h \in S^h$ denotes the finite element approximation to u . Then, there will be a balance between the error $w - u$ which measures how closely the surface fits the constraints and the error $u - u^h$, due to minimizing over a finite element subspace [Strang and Fix, 1973, pp. 132–133]. Analyzing this balance, Babuska [1973] determined that the optimal value for β is dependent on h and is given by $\beta_h = \gamma h^{-k}$, where γ is a constant and k is the degree of the complete polynomial contained in S^h . Therefore, for our quadratic finite elements, $k = 2$, and the best value for β at level j of the multi-level algorithm is $\beta_{h_j} = \gamma/h_j^2$.

5.4. Examples of Multi-Level Surface Reconstruction

Figure 7, is a schematic diagram of the structure of the multi-level surface reconstruction

¹² w is determined by the specific relaxation scheme used, but due to the size of the support of the central computational molecule, it is approximately equal to 13.

Figure 7. The structure of the multi-level surface reconstruction algorithm.



algorithm, showing three levels of resolution. The diagram depicts the relaxation processes operating at each level, as well as the fine-to-coarse and coarse-to-fine processes which transfer information between levels. The algorithm transforms a hierarchy of sparse surface depth representations, such as might be provided through the independent stereo bandpass channels, into a hierarchy of full surface representations which constitute the full $2\frac{1}{2}$ -D sketch. The constraints for the surfaces shown in the figure are random samples from a surface which varies sinusoidally in depth. It is evident that the multiple full representations output by the algorithm, describe the sinusoidal surface over a range of scales and that the accuracy of the finest representation is maintained in the coarser ones.

In this section, a number of examples of multi-level surface reconstruction are presented. We will consider the reconstruction of surfaces from artificially-generated constraints, as well as constraints generated from natural images by an implementation of the Marr-Poggio stereo theory. The performance of the multi-level algorithm is compared to that of single-level iterative algorithms. In the examples presented, the algorithm was started from identically zero initial approximations on all the levels.

In the first sequence of figures, we present synthetic examples of surface reconstructions with the purpose of illustrating the performance of the algorithm in reconstructing quadric surfaces having zero, positive, and negative Gaussian curvatures. Constraints on each level were synthesized by sampling depth along arcs on the surface. The examples, involved four levels whose grids had dimensions $N_x^{h_1} = N_y^{h_1} = 17$, $N_x^{h_2} = N_y^{h_2} = 33$, $N_x^{h_3} = N_y^{h_3} = 65$, and $N_x^{h_4} = N_y^{h_4} = 129$, with corresponding grid spacings $h_1 = 0.8$, $h_2 = 0.4$, $h_3 = 0.2$, and $h_4 = 0.1$. The relaxation method employed was the Gauss-Seidel method of equation (5.10), and a value of 2.0 was chosen for γ , giving $\beta_{h_j} = 2.0/h_j^2$.

Figure 8 shows depth constraints whose values are consistent with a cylindrical surface viewed at four resolutions. The constraints lie along arcs of greatest curvature. Figure 9 illustrates the hierarchy of full surface descriptions reconstructed by the four-level algorithm. Since the constraints on all the levels sample the same ideal cylindrical surface, the full surface representations coincide to a high degree of accuracy. Convergence was obtained after 12.0 work units. For comparison purposes, the finest level problem was isolated from the coarser levels and the same Gauss-Seidel relaxation algorithm was applied to it. Figure 10 shows the (single-level) approximation obtained after 800 work units (i.e., iterations). It is clear that we are still very far from convergence. Although the approximation is generally smooth, it has large low-frequency error components and the approximate surface lies far below its final value between constraints which are separated by fairly large distances. As predicted by the local Fourier analysis, it is precisely such low-frequency error components that local iterative algorithms have difficulty liquidating. In fact, the following characteristic phenomenon was observed. During the initial iterations, the corrections made to the approximation decreased rapidly, so that by the 800th iteration they are minute, even though

Figure 8. Constraints at four scales consistent with a cylindrical surface.




Figure 9. Hierarchy of full surface descriptions generated by the multi-level algorithm.





Figure 10. Single level approximation after 800 work units.



the error norm is still very large. Since there are 17,361 nodes in the grid, it may take on the order of $(17,361)^2$ work units to obtain the solution without the help of the coarser levels. Thus, the multi-level algorithm is vastly superior when the constraints are far apart.

Figures 11 and 12 show a synthetic example of the reconstruction of a hemispherical surface from constraints which form latitudinal circles. The hierarchy of full surface representations was obtained after 4.25 work units. Figures 13 and 14 illustrate an analogous example involving a hyperbolic paraboloid (saddle surface), where the constraints form parallel parabolic arcs. Convergence was achieved after only 2.5 work units (only a single iteration was performed on the finest level). Single level algorithms applied to the above surfaces exhibited poor convergence properties similar to the case of the cylinder.

The above examples simulate a visual situation where the surface in the scene has reflectance changes in the form of widely-spaced rulings but is otherwise free of intensity changes. This is an unlikely situation in view of the fact that the visual world is full of textures, which often arise from surface material and pigment changes. Such textures generally result in relatively densely-spaced zero-crossings forming, to a certain extent, random patterns. In turn, these zero-crossings give rise to constraints having similar properties. Figure 15 illustrates a simulation of this situation using a surface varying sinusoidally in depth. A three-level surface reconstruction algorithm was used. The constraints input on each level were 30%-density randomly-located samples of the surface depth. In addition, to simulate the effects of possible inaccuracies in the constraint values, each

Figure 11. Constraints at four scales consistent with a hemispherical surface.




Figure 12. Hierarchy of full surface descriptions generated by the multi-level algorithm.




Figure 13. Constraints at four scales consistent with a hyperbolic paraboloid.




Figure 14. Hierarchy of full surface descriptions generated by the multi-level algorithm.




Figure 15. Surface reconstruction from randomly-placed depth constraints.





Figure 16. Single level approximation after 19 work units.



sample was corrupted by zero-mean, uniformly-distributed, additive noise whose magnitude was one tenth the sample value. The algorithm generated full surface representation hierarchy in 18.75 work units. Evidently, our spring model for the influence of the constraints, with $\beta_{h_j} = 2/h_j^2$, is adequate for this case in that the additive noise has not adversely affected the quality of the reconstructed surfaces. The results after 19 work units (iterations) with single-level Gauss-Seidel relaxation on the isolated finest grid are shown in Figure 16 for comparison. It is evident that the approximation is still far from the true solution. In fact, a total of 71 work units was required to reduce the error norm to the magnitude obtained after 18.75 work units by the three-level algorithm. The saving in computation is less in this example than in the ones above because, first, only three levels were used and, second, the density of the constraints is greater. In general, the greater the density of the known depth values, the tighter the surface is constrained, and the convergence is expected to be faster. Another way to think of this is that as the average distance between constraints decreases, the efficiency of relaxation in liquidating the low-frequency Fourier components in the error increases and, therefore, the relative advantage of the multi-level algorithm is correspondingly diminished.

The next examples illustrate the performance of the multi-level algorithm using disparity constraints generated by Grimson's implementation of the Marr-Poggio stereo algorithm [Grimson, 1981a, 1981b], which includes some of the modifications suggested by Mayhew and Frisby [1980, 1981] for exploiting disparity gradient constraints along zero-crossing contours. The stereo

Figure 17. Stereo images on which the multi-level reconstruction algorithm was tested.





Figure 19. Approximation on the isolated finest grid after 950 iterations.



algorithm was run on three stereo pairs of images, shown in Figure 17, which were digitized to 320×320 pixels in 256 grey levels. The pairs from top to bottom are A) a synthesized random dot stereogram of a "wedding cake" of stacked planes, B) natural images of a portion of a coffee jar sprayed with "random dot" paint, and C) an aerial view of a set of buildings on the University of British Columbia campus. A three-channel version of the stereo algorithm was used. The resulting sparse disparity representations were reduced spatially by a factor of two, and input to a three-level surface reconstruction algorithm whose grids had dimensions $N_x^{h_1} = N_y^{h_1} = 41$, $N_x^{h_2} = N_y^{h_2} = 81$, $N_x^{h_3} = N_y^{h_3} = 161$, with corresponding grid spacings $h_1 = 0.4$, $h_2 = 0.2$, and $h_3 = 0.1$. The surface reconstructions in the examples is based on the raw disparities whose relation to depth is through a nonlinear transformation. Consequently the shapes of the reconstructed surfaces are distorted to a certain extent.

The sparse disparity constraints provided by the stereo algorithm and the hierarchy of full surface representations generated by the three-level reconstruction algorithm for the "wedding cake" stereogram are shown in Figure 18. The value $\gamma = 0.5$ was used in the algorithm, and the representations shown were generated after 16.75 work units. The three-dimensional structure of the planar surfaces is clearly evident at the three resolutions. The results can be compared to Figure 19 which shows the approximation obtained by a single-grid algorithm on the finest grid, after more than 900 work units.

Figure 20 illustrates the sparse constraints and the full surface representations obtained

Figure 18. Disparity constraints and full surface representations of "wedding cake".




Figure 20. Disparity constraints and full surface representations of the coffee jar.








Figure 21. Disparity constraints and surfaces reconstructed from the aerial view.



(with $\gamma = 0.1$) from the images of the coffee jar.¹³ The reconstruction required 16.0 work units. Finally, Figure 21 shows the sparse constraints and the reconstruction obtained from the aerial view (with $\gamma = 8.0$), after 21.875 work units. The representations are displayed as grey level images, in which darkness is proportional to disparity.

It should be noted from the above examples that the multi-level surface reconstruction algorithm, in its present form, attempts to reconstruct a single surface over the entire grid. As a consequence, serious problems arise near sharp changes in depth such as those due to partial occlusions of surfaces in the scene. The reconstructed surface gives the undesirable impression of a tablecloth thrown over a 3-D model of the scene. The source of this difficulty is discussed further in the concluding chapter, where possible ways of overcoming it are suggested.

6. GENERALIZED INTERPOLATION PROBLEMS IN VISION

The key to the solution of many problems in early vision is the imposition of constraints based upon assumptions about the visual world which are almost always true. A common assumption is that matter is cohesive; i.e., that surfaces are continuous over most of the scene. This assumption is usually introduced in the form of *smoothness constraints*, such as those characterizing the most consistent surface in visual surface reconstruction. From our study of this problem, we have seen that it is appropriate to formulate smoothness constraints within variational principles. In this chapter, we study a general class of variational principles, and we propose that the functionals characterizing these variational principles are appropriate semi-norms for formulating smoothness constraints because they possess several invariance properties which become important in applications to early vision. In order to simplify the discussion, the analysis will be in terms of interpolatory constraints and domains of infinite extent. Our visual surface reconstruction problem will be shown to be a special case of this generalized, optimal interpolation problem which is a natural generalization of the familiar curve-fitting problem involving splines.

The classical spline problem involves the minimization of the quadratic functional

$$|v|_m^2 = \int_a^b \left| \frac{d^m v(x)}{dx^m} \right|^2 dx \quad (1)$$

under the interpolatory constraints $v(x_i) = c_i$, $1 \leq i \leq N_c$ with $N_c \geq m \geq 1$, where the x_i are given distinct points in $[a, b]$ and the c_i are given real scalars. The natural setting for this problem is a vector space V formed by the class of functions whose (distributional) derivatives up to order m are in $L_2(a, b)$; that is, the class of functions which are elements of the Sobolev space of order m over $[a, b]$, $\mathcal{H}^m([a, b])$, defined in Section A.1. $|\cdot|_m$ is a semi-norm which is derived from a semi-inner product and, equipped with it, V becomes a *semi-Hilbert space*. The

¹³In this example, the constraints for the coarser channels were generated by averaging down the finest-channel disparities.

conditions imposed on the constraints ensure the existence of a unique solution $u \in S$, where S is the convex subset of V whose elements interpolate the constraints. The characterization of u as an odd-degree polynomial spline and various intrinsic properties such as the minimum norm and best approximation properties follow from the orthogonal projection theorem [Ahlberg *et al.*, 1967] (see also the proof of Theorem A.1).

Duchon [1976, 1977] and Meinguet [1979a, 1979b] describe an n -dimensional generalization of the optimal, univariate spline interpolation problem. The *generalized optimal interpolation problem* involves the minimization of the functional $|\cdot|_m^2$, where

$$|v|_m = \left(\sum_{i_1, \dots, i_m=1}^n \int_{\mathbb{R}^n} \left| \frac{\partial^m v(\mathbf{x})}{\partial x_{i_1} \dots \partial x_{i_m}} \right|^2 d\mathbf{x} \right)^{\frac{1}{2}}, \quad (2)$$

and m and n are given positive integers. The generalized interpolation problem is naturally set in a space V of functions which are elements of the Sobolev space $\mathcal{H}^m(\mathbb{R}^n)$.¹⁴ $|\cdot|_m$ is a semi-norm whose *null-space*¹⁵ is the $M = \binom{n+m-1}{n}$ dimensional space of all polynomials over \mathbb{R}^n of degree less than or equal to $m-1$: $\mathcal{N} \equiv \Pi^{m-1}(\mathbb{R}^n)$ [Schwartz, 1966, pg. 60]. Equipped with the semi-inner product corresponding to $|\cdot|_m$, V becomes a semi-Hilbert space.¹⁶

Let the finite set of distinct constraints $C = \{(\mathbf{x}_i, c_i) \mid 1 \leq i \leq N_c, \mathbf{x}_i \in \mathbb{R}^n, c_i \in \mathbb{R}\}$ contain a subset of M members such that there exists a unique element $p \in \Pi^{m-1}(\mathbb{R}^n)$ of the null space of $|\cdot|_m$ which interpolates the M constraints in the subset; that is, such that there exists a unique polynomial of degree $m-1$ which satisfies the conditions $p(\mathbf{x}_j) = c_j$ for each j which indexes a constraint of the above subset. We call such a subset an \mathcal{N} -*unisolvent subset*.

We can pose the generalized optimal interpolation problem in the following way. Given a set of constraints C which contain an \mathcal{N} -unisolvent subset, find that element $u \in S$ such that

$$|u|_m^2 = \inf_{v \in S} |v|_m^2,$$

where once again S is the set of functions in V which interpolate the constraints. The problem is well-posed because we are minimizing a semi-norm within a convex subset S of a semi-Hilbert space and, furthermore, the existence of an \mathcal{N} -unisolvent subset of constraints reduces the null-space of the functional to at most a single nonzero element of S . A solution u is then guaranteed to exist and be unique by the orthogonal projection theorem (see the proof of theorem A.1).¹⁷

¹⁴More precisely, the space V is the *Beppo Levi space* [Duchon, 1977; Meinguet, 1979a, 1979b] of order m over \mathbb{R}^n defined by $BL^m(\mathbb{R}^n) = \{v \mid \partial^\alpha v \in L_2 \text{ for } |\alpha| = m\}$, where $\alpha = (\alpha_1, \dots, \alpha_n)$ is a "vector" of positive integers and $|\alpha| = \alpha_1 + \dots + \alpha_n$; that is, it is the vector space of functions for which all the partial derivatives of (total) order m are square integrable in \mathbb{R}^n . The Beppo Levi spaces are related to the Sobolev Spaces.

¹⁵The null space of a semi-norm is the space of functions which the semi-norm maps to zero.

¹⁶According to the Sobolev inequality (see section A.1), when $m > n/2$, V is a semi-Hilbert function space of *continuous functions* [Meinguet, 1979a, 1979b].

¹⁷Moreover, as a consequence of the Sobolev inequality given in Section A.1, u will be continuous if $m > n/2$.

Why is the class of semi-norms defined in (2) important in the context of vision? As was argued recently by Brady and Horn [1981], many processes in early vision are approximately isotropic and, therefore, it seems that operators which model these processes ought to be rotationally symmetric. An example of such an operator is the $\nabla^2 G$ edge operator proposed for computing the primal sketch [Marr and Hildreth, 1980; Hildreth, 1980]. The class of semi-norms defined in (2) is of interest, since all its members $|v|_m$ are invariant under rotation and translation transformations and, moreover, if a dilation or contraction $\mathbf{x} \mapsto \lambda \mathbf{x}$ is applied to v , they are multiplied by some power of $|\lambda|$ [Duchon, 1977, pg. 86]. Therefore, corresponding interpolation methods will commute with any similarity transformations applied to the constraints. Clearly, these properties are essential for interpolation processes which contribute to the generation of the $2\frac{1}{2}$ -D sketch — the surfaces generated by surface reconstruction algorithms should not change shape as the objects in the scene undergo translations or rotations parallel to the image plane, or undergo displacements directly towards or away from the viewer.

For certain instances of $m \geq 1$ and $n \geq 1$, the general interpolation problem has familiar physical interpretations which are most often encountered in a differential form through the associated Euler-Lagrange equations. Consider first the one-dimensional case, $n = 1$. The generalized interpolation problem then reduces to the common univariate spline problem of equation (1). The particular value chosen for m determines the order of continuity of the optimal curve — as m increases, the smoothness of the solutions increases. In particular, for the case $m = 1$, $|v|_1^2 = \int_{\mathbb{R}} v_x^2 dx$ measures the energy in a string of infinite extent, and leads to interpolants having C^0 continuity. The associated Euler-Lagrange equation is $u_{xx} = 0$. C^1 continuity may be imposed on the interpolant by choosing $m = 2$. In this case, $|v|_2^2 = \int_{\mathbb{R}} v_{xx}^2 dx$ measures the strain energy of bending in a thin beam of infinite extent, and the Euler-Lagrange equation is $u_{xxxx} = 0$. This class of univariate semi-norms seems to be appropriate for imposing continuity constraints in the computation of optical flow along zero-crossing contours in the primal sketch [E.C. Hildreth, personal communication].

Next consider the generalized interpolation problem in two dimensions. For $n = 2$, the semi-norms become

$$|v|_m^2 = \int \int_{\mathbb{R}^2} \sum_{i=0}^m \binom{m}{i} \left(\frac{\partial^m v}{\partial x^i \partial y^{m-i}} \right)^2 dx dy,$$

m , once again, determining the degree of smoothness of the solution. For $m = 1$, $|v|_1^2 = \int \int_{\mathbb{R}^2} (v_x^2 + v_y^2) dx dy$ measures the potential energy related to the statics of a membrane (rubber sheet), and the associated Euler-Lagrange can be shown to be Laplace's equation, $\Delta u = 0$ [Courant and Hilbert, 1953, pg. 247]. A semi-norm of this order implicitly imposes the smoothness constraints in algorithms proposed for computing lightness [Horn, 1974], shape from shading [Ikeuchi and Horn, 1981], optical flow [Horn and Schunck, 1981], photometric stereo [Ikeuchi, 1981], etc. With $m = 2$, the smoothness of the interpolating surface is increased to C^1 , the

functional taking the form $|v|_2^2 = \int \int_{\mathbb{R}^2} (v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) dx dy$. This will be recognized as being our familiar functional representing the strain energy of the thin plate \mathcal{E}_1 with Poisson constant $\sigma = 0$ (refer to equation (2.1)), whose Euler-Lagrange equation was shown to be the biharmonic equation, $\Delta^2 u = 0$.¹⁸ As we have demonstrated, this order of smoothness seems to be most appropriate in visual surface reconstruction from, e.g., stereo information (refer also to the discussion on the “quadratic variation” in [Grimson, 1981a]).

It becomes clear that we are dealing with a class of quadratic variational problems, the order of whose Euler equations is determined by the degree of smoothness demanded of the solutions. For $m = 1$ we obtain Laplace’s equation in n dimensions, for $m = 2$ the biharmonic equation, and so on. In general, the Euler equation is an n -dimensional, linear, elliptic partial differential equation of order $2m$. Moreover, the general interpolation problems have straightforward formulations as analogous approximation problems. For example, we can define appropriate constraint terms analogous to the term \mathcal{E}_5 (equation (2.2)) for our surface approximation problem. Hence, there exists a general framework in which to solve functional approximation problems, of the type arising naturally when imposing smoothness constraints in early vision. Meaningful, problems can be formulated in any number of dimensions, and the degree of smoothness that the solutions should possess can be specified *a priori*. In this sense then, the Sobolev spaces can be viewed as ingenious formalizations of the notion of the “degree of smoothness” of admissible functions and therefore are ideal domains in which to pose and solve these problems. By specifying the (order of) the Sobolev space to which the solution should belong, we designate its position in the wide spectrum from very smooth functions to singular distributions. Satisfaction of the requirements, that the admissible space be a semi-Hilbert space and that the constraints include an \mathcal{N} -unisolvant subset will guarantee uniqueness. Needless to say, the theory of the finite element method is applicable to either the interpolation or approximation formulations, and it will dictate appropriate finite element discretization schemes for the associated variational principles.

When solving these variational principles using local, iterative algorithms such as the ones described in this paper, smoothness constraints are imposed globally over retinocentric representations by a process of *constraint propagation*. Inspired by the work of Waltz [1975], a class of algorithms called *relaxation labeling algorithms* were introduced as cooperative, constraint propagation techniques in vision and image processing by Rosenfeld, Hummel, and Zucker [1976]. Although they have seen extensive use [Davis and Rosenfeld, 1981], their generality has made them difficult to understand and, unlike the techniques and algorithms which are the subject of this paper, the foundations of most relaxation labeling schemes are unfortunately poorly-developed mathematically.

¹⁸Duchon [1977] understandably refers to the solutions as *thin plate splines* which also reflects the fact that they are natural two-dimensional generalizations of commonly-used univariate splines. In the engineering literature they are called *surface splines* [Harder and Desmarais, 1972].

Recently, some theoretical understanding has been achieved by viewing relaxation algorithms as techniques for solving constrained optimization problems (see, e.g., [Ullman, 1979b], [Faugeras and Berthod, 1981], [Hummel and Zucker, 1980]). From this new point of view, the relationships between relaxation labeling techniques and iterative solution of finite element equations arising from variational formulations become clearer — relaxation labeling schemes can be viewed as iterative algorithms for solving optimal approximation problems over closed convex *subsets* (of possible labelings) [Hummel and Zucker, 1980]. Necessary conditions for solutions (fixed points) are then expressed as sets of variational inequalities [Ciarlet, 1978; Kinderlehrer and Stampacchia, 1980] and appropriate updating rules are natural generalizations of the classical local iterative methods for solving large systems of linear [Young, 1971] or nonlinear [Ortega and Rheinboldt, 1970] equations. Moreover, if the compatibility functions among neighboring nodes are symmetric, then there exist associated variational principles defining equivalent formulations as minimization problems. Fortunately, it is possible to apply the finite element method to nonlinear problems stemming from variational inequalities [Ciarlet, 1978]. In a certain sense then, finite elements can be viewed as systematically-derived, physically-based compatibility relationships among neighboring nodes. In view of the relationships between the two techniques, it is hoped that aspects of our multi-level approach to solving the discrete finite element equations for the surface reconstruction problem may contribute to the theory of hierarchical relaxation labeling [Davis and Rosenfeld, 1981; Zucker, 1978; Zucker and Mohammed, 1979].

7. SUMMARY AND EXTENSIONS

Information about the shapes of visual surfaces that is inferred from the retinal images in the early computational stages in vision is sparse. Nevertheless, our perception is that of full (piecewise) continuous surfaces. In this paper we have proposed a hierarchical approach to the reconstruction of full surface representations, consistent with our perception of the visual world. The foundations of our paradigm are embedded in a tight mathematical formalism which at the same time seems sufficiently general to encompass many aspects of the complex information processing task which is the generation of the full $2\frac{1}{2}$ -D sketch.

Visual surface reconstruction was formulated as an optimal approximation problem having an intuitively simple physical interpretation — a thin flexible plate which is allowed to achieve an energy-minimizing state of stable equilibrium under the influence of externally-imposed constraints. This physical model led directly to an analysis in terms of the calculus of variations, and a proof that the problem is well-posed in practice. The model also suggested a class of techniques for optimally approximating the continuous solution by an equivalent discrete problem which is amenable to computational solution. We chose to apply the finite element method for reasons which include its generality, the availability of a tight theory governing its use, the simple discrete problem to which it gives rise, and its promise in vision as a systematic methodology

for constructing local representations of surfaces.¹⁹ At each step, the underlying mathematical theory assured us that, ultimately, our problem would have a unique solution that, in principle, could be computed by biologically-based mechanisms. Our search for efficient algorithms and our insights into the multi-level structure of the early processing stages in vision led us to a multi-level algorithm which solves simultaneously a hierarchy of surface approximation problems spanning a range of resolutions. The local-support processes comprising the algorithm include iterative intra-level relaxation processes, and inter-level processes which serve to communicate information between levels. The inter-level process include injections from fine grids to coarse and polynomial interpolations from coarse grids to fine. Tests on stereo data verified that our multi-level surface reconstruction algorithm meets theoretical expectations of increased speed and, moreover, generates a potentially useful hierarchy of consistent surface representations. Finally, we examined our basic surface approximation problem in a more general setting, and related it to a broader class of optimal approximation problems based on semi-norms that commute with similarity transformations applied to the constraints, a property which is important in the context of vision.

Although we have laid down the foundations of our approach primarily in terms of stereopsis, the methodology is by no means limited to the type of information produced by this particular module. Indeed, our point of view speaks to the broader issue of how to combine the information about the shapes of visible surfaces generated by various vision modules into a self-consistent whole. Several possibilities arise, some of which we will now consider briefly.

The simultaneous assimilation of information from different sources can be realized by defining more sophisticated penalty functions to replace \mathcal{E}_5 in Chapter 2. For example, in the case of depth constraints from, say, stereo and motion, we can straightforwardly introduce additional terms of the same form as \mathcal{E}_5 for each process. In terms of our plate model, we introduce two sets of imaginary pins with attached springs, and allow the possibility of a constraint generated by stereo to coincide with one provided by motion. Imperfections in the retinal images are likely to affect the two processes in different ways, and moreover each will in its own way sporadically fall prey to gross misinterpretations of the information in the primal sketches. Whatever the situation, our physical model assumes an energy-minimizing state, and the resulting surface is an optimal compromise in view of the constraints provided. In places where the information is consistent, the final interpretation is reinforced. In places where there is a conflict, it is resolved by competition with nearby constraints from both processes.

The influence of each constraint may be controlled, possibly dynamically by the processes themselves, by assigning different values to each spring constant. For example, different *confidence values* may be given to individual constraints generated by the stereo matcher, according to

¹⁹On this latter point we should mention again that the finite element method allows us to handle domains of complex shape, natural boundary conditions, and to set up nonuniform discretizations of the domain — e.g., to vary the resolution across the domain.


regional statistics of the rate of successfully matched zero-crossings, which may have to be computed anyway [Marr and Poggio, 1979; Grimson, 1981a, sections 2.5 and 3.4]. The extent of the constraint's influence on the surface may also be varied by extending our model to that of an inhomogeneous plate, whose flexibility varies over the domain. Numerous possibilities exist for defining weighting functions to apply to the strain energy density of the plate. All such proposals for modifying the form of the functional must be first be shown to lead to well-posed problems, by extensions of the analysis carried out in Chapters 3 and 4.

Another important issue is how to incorporate information other than measurements of the distance to the surface. An important class of processes generate cues about visual surfaces in the form of *local orientation measurements*. Examples in this class are the analysis of occluding contours [Marr, 1977], as well as "shape-from" processes such as shape from shading [Horn, 1975], contours and texture [Kender, 1980; Stevens, 1981; Witkin, 1981], regular patterns [Kanade, 1981], etc. The finite element method provides a general way of handling orientation constraints through the use of elements with degrees of freedom which include the first partial derivatives of the surface. An appealing example is Adini's rectangle which was described in Section 4.1. Surface representations based on this element would make explicit the information about the local slope of surfaces, as well as their distance from the viewer. Discrete problems derived by applying this element would correspond to a coupled system of two discrete Euler-Lagrange equations for the plate, a fourth-order equation for the displacements at the nodes, and a second-order equation for the slopes. On the other hand, we pay a price for this added capability — the dimension of the finite element space is tripled, making the resulting discrete system even larger. Nevertheless, the price may turn out to be worth paying in order to obtain useful surface representations and, moreover, it may not be too high when massively parallel computational mechanisms are contemplated.

A different way of handling orientation constraints which can be used with our simple quadratic elements is, once again, by the use of appropriate penalty functions. In terms of our model, we can imagine the situation for a single constraint and a surface patch as illustrated in Figure 22. Here, we attach a spring between the surface normal, an imaginary quill rigidly fixed to the plate's surface at a particular point, and the orientation constraint, another quill emanating from the same point, but having a fixed orientation in space. Given this arrangement, the surface is "pulled" locally so that its orientation tends to align with that of the constraint. The appropriate penalty term is the potential energy in the spring. This energy can be expressed straightforwardly in terms of the first partial derivatives of the quadratic surface patch within the element, and ultimately in terms of the node displacements.

A more immediately important issue, one that was raised in reference to the examples of surface reconstruction presented earlier, is that of dealing with *depth discontinuities*. In its present form, the surface approximation algorithm can deal in a meaningful way with scenes containing

Figure 22. Physical model of the effect of an orientation constraint.



only a single surface. This is due to the fact that it does not incorporate the notion of an *occluding contour*; that is to say, it attempts to fit a single surface over the whole sparse $2\frac{1}{2}$ -D sketch, interpolating indiscriminately across contours which correspond to places where surfaces in the scene occlude one other from the viewer. Clearly, this action is inappropriate since the surfaces on either side of the occluding boundary ought to have no influence on one another. Moreover, the variational principle for surface approximation was based on the *small deflection* theory of the plate²⁰ and, consequently, we expect our surface to behave strangely in the vicinity of a large change in depth, resulting in, for example, a Gibb's phenomenon similar to that observed when approximating discontinuous functions with Fourier series.

How can these depth discontinuities be detected and how do we prevent interpolation across them? Grimson [1981a, section 9.4] noted the importance of this question and made some speculations about possible answers to the first of its two parts. The feasibility of his suggestions remains an open question. Here, we would like to propose another approach that is again suggested by our physical model. In places of sharp changes in depth (or surface orientation), the strain energy in the plate will be locally high. Measuring this energy locally is a simple matter — we use the energy inner product to compute the strain energy norm $a_h(v^h, v^h)^{\frac{1}{2}}$ over the element domains. Points of high strain energy are likely candidates for inferring the presence of discontinuities in

²⁰The large deflection plate bending theory is considerably more complicated and leads to an Euler-Lagrange equation in the form of two coupled, nonlinear, fourth-order partial differential equations known as *von Karmann's equations* (see, e.g., [Landau and Lifshitz, 1970] or [Mansfield, 1964])

depth. We can also exploit our expectations about the world for added constraint, and assert that, since most of the retinal image is made up of coherent surfaces, occlusions in depth are likely to form contours in the image and not be sparsely-occurring points. Hence, we look for *contours* along the surface of the plate, where the strain energy is high. Having located the occluding contours, the answer to the second part of the question is simple, in principle. To prevent interpolation across different surfaces, we "break" the plate along occluding contours. Mathematically speaking, this is done by removing plate elements along the occluding contour, thereby introducing free boundaries.

Our multi-level approach to surface reconstruction constitutes a computational paradigm which has contributed toward a more complete understanding of the generation of the $2\frac{1}{2}$ -D sketch. Many details such as the combination of information from the various modules in early vision and the isolation of depth discontinuities remain to be worked out rigorously within the paradigm. In addition, a number of exciting issues are raised. For example, how can the hierarchy of surface representations generated by the algorithm be used to advantage during later computational stages in which three-dimensional, object-centered representations are generated and objects are recognized. Similar implications directed to the related field of robotics and manipulation also suggest themselves. Research addressing some of these issues is currently in progress.

ACKNOWLEDGEMENTS

I would like to thank Shimon Ullman for pointing out the literature on multi-grid methods and suggesting the possibility of their application to visual surface interpolation. Mike Brady, Eric Grimson, Tomaso Poggio, Whitman Richards, and Shimon Ullman contributed to the work through numerous valuable discussions. In addition, Eric Grimson generously provided technical assistance with his stereo vision software for the LISP Machine. Drafts of this paper were read by Mike Brady, Eric Grimson, Ellen Hildreth, Berthold Horn, and Shimon Ullman. Their comments are appreciated greatly.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505 and in part by National Science Foundation Grant 79-23110MCS. The author was supported by postgraduate scholarships from the Natural Sciences and Engineering Research Council of Canada and from Fonds F.C.A.C. pour l'aide et le soutien à la recherche, Québec, Canada.

A. THE FINITE ELEMENT METHOD

When it is impossible to derive an analytical solution to a continuous variational principle, it is usual to attempt an approximation by defining a *discrete problem* which is similar to the continuous one and which leads to a *discrete solution*. To this end, we will first state an *abstract variational principle* which will lead us to an optimal approximation to the exact solution. The variational principle is called *abstract* inasmuch as it represents a formulation which is common to a variety of physical problems, such as the physical model for our surface reconstruction problem. We will also state theorems which give conditions guaranteeing the existence and uniqueness of the approximate solution and, in addition, we will discuss the optimal properties of the proposed approximation.

The abstract variational principle and the associated theorems are stated in a form which is convenient for the application of the *finite element method*, a powerful technique for obtaining, by numerical means, discrete solutions to variational problems.¹ The following sections develop the mathematical machinery which we will require to successfully apply the method. Key mathematical ideas include a set of Hilbert spaces (the Sobolev spaces) and their norms, a bilinear form (the energy inner product) which is naturally associated with the specific problem, and certain optimal properties of the (Ritz) approximation over finite dimensional subspaces. These ideas lead to a clean and precise theory governing the application of the finite element method, even for complicated geometries. Comparatively tight theories are unavailable for alternate approximation techniques which naturally arise from nonvariational problem statements; e.g., the finite difference method which can be applied to equivalent formulations in terms of differential operator equations (such as Euler-Lagrange equations). Excellent accounts of the mathematical theory of the finite element method are [Ciarlet, 1978], [Oden and Reddy, 1976], and [Strang and Fix, 1973]. An extensive development from an engineering point of view is presented in [Zienkiewicz, 1977].

A.1. The Sobolev Spaces

Fundamental to finite element analysis are a set of spaces called the *Sobolev spaces* (see, e.g., [Agmon, 1965], [Adams, 1975]). They are a generalization of the familiar L_2 space which consists of all functions $v: \Omega \subset \mathbb{R}^n \mapsto \mathbb{R}$ (where Ω is a bounded domain) whose L_2 norm over Ω

$$\|v\|_{L_2, \Omega} = \left(\int_{\Omega} |v(\mathbf{x})|^2 d\mathbf{x} \right)^{\frac{1}{2}}$$

is finite. We denote the partial derivatives of v by the notation

$$\partial^{\alpha} v(\mathbf{x}) = \left(\frac{\partial}{\partial x_1} \right)^{\alpha_1} \cdots \left(\frac{\partial}{\partial x_n} \right)^{\alpha_n} v(\mathbf{x}),$$

¹The finite element method was conceived by Courant [1943].

where $\alpha = (\alpha_1, \dots, \alpha_n)$ is a multi-index of positive integers α_i .² The *Sobolev norm* of order m over Ω combines the L_2 norms of all partial derivatives of v up to order m :

$$\|v\|_{m,\Omega} = \left(\sum_{|\alpha| \leq m} \int_{\Omega} |\partial^{\alpha} v|^2 d\mathbf{x} \right)^{\frac{1}{2}}, \quad (1)$$

where $|\alpha| = \alpha_1 + \dots + \alpha_n$. The Sobolev space of order m over Ω is then defined by

$$\mathcal{H}^m(\Omega) = \{v \mid \|v\|_{m,\Omega} < \infty\}. \quad (2)$$

Clearly, $\mathcal{H}^0 \equiv L_2$.

We will also require the associated semi-norm

$$|v|_{m,\Omega} = \left(\sum_{|\alpha|=m} \int_{\Omega} |\partial^{\alpha} v|^2 d\mathbf{x} \right)^{\frac{1}{2}} \quad (3)$$

which includes only the derivatives of order m exactly. It is a semi-norm because it is zero if $v = \pi_{m-1} \in \Pi^{m-1}(\Omega)$, where Π^{m-1} is the space of polynomials of degree $m-1$ defined in Ω [Strang and Fix, 1973, pg. 298].

Since the Sobolev norms are sums of L_2 norms, they have associated inner products and, therefore, the Sobolev spaces are *Hilbert spaces*.³ Let $C^q(\Omega)$ denote the class of functions which have continuous partial derivatives of all orders up to order q . A fundamental embedding property of the Sobolev spaces is given by the *Sobolev inequality* which states that $C^q \subset \mathcal{H}^m$ if and only if $m - q > n/2$, where n is the dimension of \mathbb{R}^n .

A.2. An Abstract Variational Principle

Let V be a normed vector space with norm $\|\cdot\|$, and S be a nonempty subset of V . Moreover, let $a(\cdot, \cdot): V \times V \mapsto \mathbb{R}$ be a continuous bilinear form and $f: V \mapsto \mathbb{R}$ be a continuous linear form.

Definition 1 — abstract (quadratic) variational principle

The problem: find an element u^S such that

$$u^S \in S \quad \text{and} \quad \mathcal{E}(u^S) = \inf_{v^S \in S} \mathcal{E}(v^S), \quad (4)$$

where the functional $\mathcal{E}: V \mapsto \mathbb{R}$ is defined by

$$\mathcal{E}(v) = \frac{1}{2}a(v, v) - f(v), \quad (5)$$

²The derivatives are to be interpreted in the generalized (distributional) sense, but when a derivative exists in the classical sense, it is equal to the generalized derivative (see, e.g., [Schwartz, 1966]).

³Extensions of the definition of Sobolev spaces and norms have been made to negative and nonintegral order m (see, e.g., [Agmon, 1965], [Adams, 1975]).

will be referred to as the abstract variational principle.

Theorem 1 — existence and uniqueness

Assume in addition that

- (i) the space V is complete,⁴
- (ii) S is a closed convex subset of V ,
- (iii) the bilinear form $a(\cdot, \cdot)$ is symmetric,
- (iv) $a(\cdot, \cdot)$ is V -elliptic;⁵ i.e., there exists a constant $\alpha > 0$ such that

$$\forall v \in V, \quad a(v, v) \geq \alpha \|v\|^2, \quad (6)$$

then the abstract variational principle has a unique solution $u^S \in S$.⁶

Proof. [Ciarlet, 1978, pg. 3] The bilinear form $a(\cdot, \cdot)$ is an inner product over the space V . Since it is continuous, it is bounded,⁷ and because it is also V -elliptic, there exist constants α and μ such that $\alpha \|v\|^2 \leq a(v, v) \leq \mu \|v\|^2$. Therefore, the norm associated with the inner product is equivalent⁸ to the given norm $\|\cdot\|$, and V becomes a Hilbert space when it is equipped with this inner product. According to the *Reisz representation theorem* (see, e.g., [Oden, 1979] or [Yosida, 1971]) there exists an element $u \in V$ such that

$$\forall v \in V, \quad f(v) = a(u, v), \quad (7)$$

and because $a(\cdot, \cdot)$ is symmetric,

$$\mathcal{E}(v^S) = \frac{1}{2} a(v^S, v^S) - a(u, v^S) = \frac{1}{2} a(v^S - u, v^S - u) - \frac{1}{2} a(u, u).$$

Therefore, the minimum of $\mathcal{E}(v^S)$ and the minimum of $a(v^S - u, v^S - u)$ as v^S ranges over the set S are achieved by the same element $u^S \in S$. In other words, solving the abstract variational principle is equivalent to finding an element $u^S \in S$ which is closest to u with respect to the norm $a(\cdot, \cdot)^{\frac{1}{2}}$:

$$a(u - u^S, u - u^S)^{\frac{1}{2}} = \inf_{v^S \in S} a(u - v^S, u - v^S)^{\frac{1}{2}}. \quad (8)$$

By the *projection theorem*, (see, e.g., [Oden, 1979] or [Yosida, 1971]) the solution is the projection of u onto S with respect to the inner product $a(\cdot, \cdot)$, and its existence and uniqueness is assured by the fact that S is a closed convex subset of V . ■

⁴That is to say, it is a *Banach space*.

⁵ V -ellipticity means that the bilinear form is *positive definite*; i.e., $a(v, v) = 0$ if and only if $v = 0$.

⁶Theorem 1 is a generalization of the familiar theorem for the existence of a unique solution to a (quadratic) minimization problem in mathematical programming (see e.g. [Luenberger, 1973]).

⁷A bilinear form is continuous if and only if it is bounded; i.e., there exists a constant μ such that $|a(u, v)| \leq \mu \|u\| \|v\|$ [Rektorys, 1980, pg. 111].

⁸Two norms $\|\cdot\|$ and $\|\cdot\|'$ on a linear vector space V are called *equivalent* if the corresponding metrics are equivalent. This amounts to the existence of two positive constants c_1 and c_2 such that $c_1 \|v\| \leq \|v\|' \leq c_2 \|v\|$.

The abstract formulation encompasses linear variational problems which are posed classically in terms of variational principles involving the minimization of quadratic functionals $\mathcal{E}(v) = \frac{1}{2}a(v, v) - f(v)$ over an *admissible space of functions* $v \in V$. Such functionals often represent the potential energy of a physical system, $a(v, v)$ being the second-degree term which is the *strain energy* in the function v ($f(v)$ is a first-degree term). The associated inner product $a(v, w)$ is the *energy inner product* which is intrinsic to the particular variational principle, and is defined for all admissible functions v and w .

It is clear from the above discussion that the admissible space V must be complete and that $\mathcal{E}(v)$ must be well-defined for all $v \in V$ (i.e., that V must be a *space of finite energy*). The Sobolev spaces fulfill these conditions. Their use as generalized energy spaces is natural in the sense that, for a given variational principle, the energy inner product is well-defined over a Sobolev space \mathcal{H}^m , where m is the highest order of partial derivative of v which occurs in $a(v, v)$. In general then, V is the space \mathcal{H}^m whose natural norm is the Sobolev norm $\|\cdot\|_m$.

The role of the subset S is in approximating the exact solution u . Although u is usually impossible to obtain over the full admissible space V , it may be relatively straightforward to optimally approximate it by an element $u^S \in S$, especially if S is taken to be a *closed subspace* of V . The approximation is optimal in the sense of equation (8). In the ensuing discussion, we will restrict ourselves to the special case where S is a closed subspace of V . The approximate solution of the abstract variational principle may then be characterized by the following theorem.

Theorem 2 — variational equation

If S is a closed subspace of V , then $u^S \in S$ is a solution of the abstract variational principle if and only if it satisfies the variational equation

$$\forall v^S \in S, \quad a(u^S, v^S) = f(v^S). \quad (9)$$

Proof. If u^S minimizes \mathcal{E} over S , then for any ϵ and $v^S \in S$,

$$\begin{aligned} \mathcal{E}(u^S) &\leq \mathcal{E}(u^S + \epsilon v^S) = \frac{1}{2}a(u^S + \epsilon v^S, u^S + \epsilon v^S) - f(u^S + \epsilon v^S) \\ &= \mathcal{E}(u^S) + \epsilon[a(u^S, v^S) - f(v^S)] + \frac{1}{2}\epsilon^2 a(v^S, v^S). \end{aligned}$$

Therefore,

$$0 \leq \epsilon[a(u^S, v^S) - f(v^S)] + \frac{1}{2}\epsilon^2 a(v^S, v^S),$$

and since this must be true for small ϵ , both positive and negative, it follows that $a(u^S, v^S) = f(v^S)$.⁹ ■

⁹In the general case where S is not a closed subspace, but is only a closed convex subset of V as required by condition (ii) of Theorem 1, the solution u^S must satisfy the *variational inequality* $a(u^S, v^S - u^S) \geq f(v^S - u^S)$ [Ciarlet, 1978, pg. 3].

We will now discuss some important properties of the solution of the abstract variational principle. First, from the proof of the theorem, it is clear that (9) is the well-known condition for *the vanishing of the first variation* of ϵ at u^S , in the direction of v^S . In particular, if S is the whole space V , then the solution satisfies

$$a(u, v) = f(v), \quad (10)$$

and the first variation at u vanishes in every direction v . Setting $u = v$, we have $a(u, u) = f(u)$ and hence

$$\mathcal{E}(u) = \frac{1}{2}a(u, u) - f(u) = -\frac{1}{2}a(u, u);$$

i.e., *at the minimum, the strain energy is the negative of the potential energy*. Now, $\mathcal{E}(u) \leq \mathcal{E}(u^S)$ since u is minimal over a wider class of functions. Then,

$$a(u^S, u^S) \leq a(u, u),$$

and so *the strain energy in u^S always underestimates the strain energy in u* . Moreover, since u^S is the projection of u onto the subspace S , the error $e^S = u - u^S$ is orthogonal to S :

$$a(e^S, v^S) = 0 \quad \forall v^S \in S.$$

In particular, $a(e^S, u^S) = 0$ or $a(u, u^S) = a(u^S, u^S)$ and

$$a(e^S, e^S) = a(u, u) - a(u^S, u^S);$$

i.e., *the energy in the error equals the error in the energy* (the Pythagorean theorem holds).

A.3. The Ritz Approximation

The key condition in the hypothesis of Theorem 2 is that the subspace S be closed. How can we ensure this? One possibility arises from the fact that *finite dimensional subspaces are always closed*. A number of classical methods for solving variational problems, called *direct methods*, are based on this.¹⁰ One of these, the *(Raleigh-) Ritz method* [Ritz, 1908; Mikhlin, 1964; Rektorys, 1980] is of fundamental importance when a variational principle is involved. In the Ritz method, we choose a finite dimensional subspace

$$S = S^h \equiv \{v^h \mid v^h = \sum_{i=1}^N v_i \phi_i\} \quad (11)$$

where ϕ_1, \dots, ϕ_N are independent *basis functions* which span S^h and v_1, \dots, v_N are unknown real parameters.

¹⁰Direct methods include the *method of weighted residuals* whose special cases include *collocation methods* and the *Galerkin method*, the *method of orthonormal* (e.g., *Fourier*) *series*, and the *least squares method*, (see, e.g., [Finlayson, 1972], [Mikhlin, 1964], [Rektorys, 1980], [Zienkiewicz, 1977]).

By Theorem 1, the approximate solution to the variational principle is the unique element $u^h \in S^h$ which is the projection of u onto S^h . This amounts to choosing parameters v_i which satisfy the *discrete variational principle*

$$\mathcal{E}(u^h) = \inf_{v^h \in S^h} \mathcal{E}(v^h) = \inf_{v_1, \dots, v_N \in \mathbb{R}} \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N a(\phi_i, \phi_j) v_i v_j - \sum_{i=1}^N f(\phi_i) v_i, \quad (12)$$

or, by Theorem 2, the associated variational equation

$$\forall v^h \in S^h, \quad a(u^h, v^h) = f(v^h)$$

which is, in fact, a linear system of algebraic equations

$$\sum_{j=1}^N a(\phi_i, \phi_j) u_j = f(\phi_i), \quad 1 \leq i \leq N.$$

These equations can be written in the compact matrix form

$$\mathbf{A} \mathbf{u} = \mathbf{f}, \quad (13)$$

where $\mathbf{A} \in \mathbb{R}^{NN} = [a(\phi_i, \phi_j)]$ and where $\mathbf{f} \in \mathbb{R}^N = [f(\phi_i)]$ and $\mathbf{u} \in \mathbb{R}^N = [u_i]$, called the *discrete variational equation*. Since the matrix of coefficients \mathbf{A} is nonsingular, the *discrete solution* is given by $\mathbf{u} = \mathbf{A}^{-1} \mathbf{f}$, although for the problem at hand \mathbf{A} is huge, so it is usually impractical to compute \mathbf{A}^{-1} directly. In the next section, we describe special types of basis functions which ultimately lead to practical *iterative* solutions.

A.4. Finite Element Spaces

The Ritz method has given us a *discrete solution* $u^h = \sum_{i=1}^N u_i \phi_i$ which is optimal in the sense that the energy in the error, as measured in the natural energy norm $a(u - u^h, u - u^h)^{\frac{1}{2}}$, is as small as possible. In the *classical* Ritz method, the basis functions ϕ_i are generally chosen to be fairly complicated functions which have global support over the domain in question (e.g., trigonometric functions) [Mikhlin 1964; Rektorys, 1980]. Although this choice may be beneficial for analytic purposes, it renders the method unsuitable for numerical computation. The problem is overcome by the *finite element method* which is a systematic procedure for constructing finite dimensional approximating subspaces S^h , called *finite element spaces*, which are very convenient for numerical computation. In certain forms, the method may be considered to be a special instance of the Ritz method in which the basis functions are *simple functions having local support*. In the ensuing discussion, we will restrict ourselves to a domain $\bar{\Omega}$ which is a polygon in \mathbb{R}^2 with boundary $\partial\Omega$.¹¹ The following are basic characteristics of the construction in its simplest form:

- (i) A “*triangulation*” \mathcal{T}^h is established over the domain: $\bar{\Omega} = \bigcup_{E \in \mathcal{T}^h} E$; that is, the domain is partitioned into the union of subdomains $E \in \mathcal{T}^h$ called *finite elements*, such that the E are closed sets with nonempty interiors and polygonal boundaries. The elements are usually

¹¹The theory has been extended to domains with curved boundaries in any number of dimensions.

adjacently-placed triangles or rectangles which overlap only at the *inter-element boundaries*. Associated with the triangulation is its *fundamental length* h .¹²

- (ii) The elements are considered to be interconnected at a discrete number of points on the inter-element boundaries which are called the *nodes* of the triangulation. The unknown real parameters of the discrete problem are the *nodal variables*, the values of the solution (and/or possibly of its derivatives) at the nodes.
- (iii) Associated with the triangulation is a space of functions S^h defined over $\bar{\Omega}$. Defined within each element E is a finite dimensional space $P^E = \{v^h|_E \mid v^h \in S^h\}$ consisting of functions which are polynomials (or ratios of polynomials). The polynomials $p^E \in P^E$ represent a local approximation of the solution within E , and are uniquely determined by the nodal variables associated with the element.
- (iv) In certain elements, the functions p^E may have to satisfy the *essential boundary conditions* of the problem.

While the classical Ritz method is limited to geometrically simple domains $\bar{\Omega}$, in the finite element method this limitation occurs only within the element itself. Consequently, it is possible to "assemble" complicated configurations from simple element shapes. Several factors contribute to the success of the finite element method from a computational point of view. Firstly, due to the fact that in the element interiors the solution is approximated by a low-order polynomial in x and y , the computations required to compute the discrete functional in (12) or, equivalently, to compute the entries of matrix A of the discrete variational equation (13) are often simple. Secondly, it can be shown that, associated with the local polynomial functions, there exists a canonical set of basis functions ϕ_i spanning S^h which are also piecewise polynomials and which have *local support*. A will therefore be *sparse* and *banded*; that is, most of its relatively few nonzero entries will lie near the main diagonal. Thirdly, when the problem is well-posed in terms of a variational principle, $a(\cdot, \cdot)$ will be symmetric and S^h -elliptic, which guarantees that A will be *nonsingular*, *symmetric*, and *positive definite*. In addition to these clear merits, piecewise polynomials are remarkable in that they are optimal in terms of their approximation properties and in that these properties are essential for proving convergence of the method [Strang and Fix, 1973, pg. 153].

The convergence properties of the finite element method are an important issue. The object of the Ritz method is to find optimal values for the nodal variables (which are the parameters of the discrete solution) by minimizing the discrete functional $\mathcal{E}(v^h)$. This suggests immediately the possibility of approximating the exact solution u by a *minimizing sequence* of discrete solutions to discrete problems associated with a family of subspaces S^h whose fundamental length h has limit zero. Although the approximation is known by (8) to be optimal in terms of the norm $a(\cdot, \cdot)^{\frac{1}{2}}$, it is more convenient to analyze the error in terms of the natural Sobolev norm $\|\cdot\|_m$ of $V \subset \mathcal{H}^m$. The following theorem gives a sufficient condition for the convergence of such a sequence.

¹²The fundamental length h of the triangulation \mathcal{T}^h is the maximum "radius" of the elements. As the subdivision is made finer, the number of elements increase and $h \rightarrow 0$.

Theorem 3 — C  a's Lemma

Since there exists a constant C independent upon the subspace S^h such that

$$\|u - u^h\| \leq C \inf_{v^h \in S^h} \|u - v^h\|, \quad (14)$$

then a sufficient condition for convergence is that there exists a family S^h of subspaces of the space V such that, for each $u \in V$, $\lim_{h \rightarrow 0} \|u - u^h\| = 0$; i.e.,

$$\lim_{h \rightarrow 0} \inf_{v^h \in S^h} \|u - v^h\| = 0.$$

Proof. Equation (14) follows from (8) due to the continuity and V -ellipticity of $a(\cdot, \cdot)$. Moreover, $C = \sqrt{\frac{\mu}{\alpha}}$ is a constant independent upon the subspace S^h . ■

We see then that an estimation of the error reduces to finding the distance between the exact solution u and the subspace S^h — a problem in *approximation theory*. The basic hypothesis about the finite element space S^h was that the finite-dimensional space P^E within each element E is a space of polynomials. If we assume that the space contains the complete polynomials of degree k (i.e., $\Pi^k \subset P^E$), it can be shown in general that the approximation error in the s^{th} derivative, where $s \neq m$, is of the form

$$\|u - u^h\|_s = O(h^{k+1-s} + h^{2(k+1-m)}) \quad (15)$$

[Strang and Fix, 1973, pg. 107]. On the other hand, because the approximation minimizes the strain energy, the order of convergence of the m^{th} derivative is better. It is order $h^{2(k+1-m)}$.

The convergence properties implied by C  a's lemma are contingent upon the finite element spaces S^h being subspaces of the admissible space V . In view of this, if the energy inner product $a(v, v)$ involves partial derivatives of v of order m so that $V \subset \mathcal{H}^m(\bar{\Omega})$, ensuring convergence amounts to imposing the following two requirements on the local functions p^E :

- (i) **Completeness condition:** As the size of any element tends to zero, the function p^E must be such that a constant value of the m^{th} derivative will be attainable in the element subdomain; i.e. we must have $k \geq m$ so that $P^E \subset \mathcal{H}^m(E)$, $\forall E \in \mathcal{T}^h$.
- (ii) **Conformity condition:** All derivatives of p^E of order less than m must be continuous across inter-element boundaries; that is, $S^h \subset C^{m-1}(\bar{\Omega})$.

The two requirements are necessary and sufficient for $S^h \subset \mathcal{H}^m(\Omega)$ when the p^E are polynomials (or ratios of polynomials). Another way of stating the completeness condition is that the local polynomials must be able to reproduce a *state of constant strain* — any solution which is a polynomial of degree m . When the local polynomials satisfy the conformity condition, the elements are called *conforming finite elements*, and their use leads what are referred to as *conforming finite element methods*.

A.5. Nonconforming Elements

In the above discussion, we assumed that conforming finite element methods approximate the solution u of $\mathcal{E}(u) = \inf_{v \in S} \mathcal{E}(v)$ by the solution u^h of $\mathcal{E}(u^h) = \inf_{v^h \in S^h} \mathcal{E}(v^h)$ where S^h is a *subspace* of S . This is a global condition on the approximation which is often violated for reasons of computational convenience. For instance, it may be violated by dropping the element conformity condition.¹³ Elements which do so are called *nonconforming elements*. They are often used in practice for higher-order problems because conforming elements for such problems are unnecessarily complicated or must have a large number of degrees of freedom in order to satisfy the inter-element conformity conditions.

If nonconforming elements are used, it is clearly impossible to evaluate the true energy functional $\mathcal{E}(v^h)$, due to the singularities in the m^{th} derivatives of v^h which occur at the element boundaries. To avoid this problem, we can simply ignore the discontinuities between elements by computing the strain energies within each element and then summing the individual contributions; that is, for the original energy inner product $a(\cdot, \cdot)$, we substitute the *approximate energy inner product* $a_h(\cdot, \cdot)$ defined by the bilinear form

$$a_h(\cdot, \cdot) = \sum_{E \in \mathcal{T}^h} a(\cdot, \cdot)|_E, \quad (16)$$

where the notation $|_E$ means a restriction to the element domain. The *approximate variational principle* is then the problem of finding a $u^h \in S^h$ which minimizes the functional

$$\mathcal{E}_h(v^h) = \frac{1}{2} a_h(v^h, v^h) - f(v^h), \quad (17)$$

and the necessary condition for the vanishing of the first variation becomes

$$\forall v^h \in S^h, \quad a_h(u^h, v^h) = f(v^h). \quad (18)$$

Following in the spirit of the conforming case, we must determine sufficient conditions for the existence and uniqueness of the solution u^h to the approximate variational principle, as well as under what conditions this approximate solution converges to the exact solution u as $h \rightarrow 0$. The conditions are natural extensions of Theorem 1 and Céa's lemma, and are given in the following theorem.

Theorem 4 — existence and uniqueness (nonconforming case)

If

- (i) there exists a mapping $\|\cdot\|_h: S^h \rightarrow \mathbb{R}$ which is a norm over S^h ,
- (ii) $a_h(\cdot, \cdot)$ is bounded and S^h -elliptic, in that there exists a constant $\alpha_h > 0$ such that

¹³This violation is an example of a so called *variational crime* [Strang and Fix, 1973, Chapter 4]. Besides violation of the element conformity condition, variational crimes also include inexact evaluation of the functional $\mathcal{E}(v^h)$ (i.e., of the quadratic form $a(u^h, v^h)$ and linear form $f(v^h)$) such as by numerical integration, as well as various techniques for the approximation of essential boundary conditions.

$$\forall v^h \in S^h, \quad a_h(v^h, v^h) \geq \alpha_h \|v^h\|_h^2,$$

then the approximate variational principle has a unique solution $u^h \in S^h$.

Proof. Refer to the discussion in [Ciarlet, 1978, Chapter 4]. ■

On the other hand, to obtain convergence, we impose a stronger condition, *uniform S^h -ellipticity*, which requires that there exist a constant $\tilde{\alpha} > 0$, independent of h , such that

$$\forall S^h, \quad \forall v^h \in S^h, \quad a_h(v^h, v^h) \geq \tilde{\alpha} \|v^h\|_h^2. \quad (19)$$

Convergence is then guaranteed by the following theorem.

Theorem 5 — Strang's lemma

Given a family of discrete problems for which the associated approximate energy inner products are uniformly S^h -elliptic, then there exists a constant C , independent of S^h , such that

$$\|u - u^h\|_h \leq C \left(\inf_{v^h \in S^h} \|u - v^h\|_h + \sup_{w^h \in S^h} \frac{|a_h(u, w^h) - f(w^h)|}{\|w^h\|_h} \right)^{\frac{1}{2}}. \quad (20)$$

Proof. See [Ciarlet, 1978, pg. 210]. ■

Strang's lemma is a generalization of Céa's lemma for conforming elements — in addition to the usual approximation error term, we have the (inf) term which measures the *consistency error* of the nonconforming space. Since the difference $a_h(u, w^h) - f(w^h)$ is zero for all $w^h \in S^h$ when $S^h \subset V$, the consistency error for conforming spaces is identically zero and Strang's lemma reduces to Céa's lemma. However, for the nonconforming case, convergence is obtained if the *consistency condition*,

$$\lim_{h \rightarrow 0} \sup_{w^h \in S^h} \frac{|a_h(u, w^h) - f(w^h)|}{\|w^h\|_h} = 0, \quad \forall v^h \in S^h, \quad (21)$$

is satisfied.

The consistency condition was first recognized empirically and was stated in the form of a simple test known as the *patch test*. Subsequently, Strang proved mathematically that the patch test was indeed a test for consistency and, by essentially making the above arguments, that nonconforming elements which pass it will yield convergence (see [Strang and Fix, 1973, Chapter 4]). The test remains a convenient one to apply.

Theorem 6 — the patch test

Suppose that the energy inner product $a_h(u, v)$ contains derivatives of order at most m and the nonconforming space S^h contains all polynomials π_m of degree at most m . If the nonconforming finite element method recovers all solutions which are polynomials of degree at most m , then the patch test is passed, and $\lim_{h \rightarrow 0} \|u - u^h\|_h = 0$.

In other words, suppose that we put an arbitrary patch of nonconforming elements associated with the nonconforming space S^h in a state of constant strain; that is, we impose $v^h = \pi_m \in \Pi^m$ on the displacements at nodes around the patch boundary. Because the completeness condition of the previous section is still binding on nonconforming elements, this polynomial is both an element of S^h and an element of V , hence its consistency error is zero. The conforming (Ritz) solution to (12) or (13) and the nonconforming, discrete solution of the approximate variational principle ought then to be identical and equal to π_m . The test is to determine whether this is indeed the case.

B. ITERATIVE SOLUTION OF LARGE LINEAR SYSTEMS

The approximation of a variational principle by direct methods such as the finite element method (or the approximation of a boundary value problem by finite differences) gives rise to a system of simultaneous algebraic equations. For quadratic functionals (or linear boundary value problems), the system will be linear. In this chapter, we consider the problem of solving, by iterative means, systems of linear equations of the form

$$\sum_{j=1}^N a_{ij} u_j = f_i, \quad 1 \leq i \leq N, \quad (1)$$

where the coefficients a_{ij} and the values f_i are known. The system may also be written as the matrix equation

$$\mathbf{A}\mathbf{u} = \mathbf{f}, \quad (2)$$

where $\mathbf{A} \in \mathbb{R}^{NN} = [a_{ij}]$ is a nonsingular matrix and where $\mathbf{f} \in \mathbb{R}^N = [f_i]$ is a column vector. We wish to solve the system for the column vector $\mathbf{u} \in \mathbb{R}^N = [u_i] = \mathbf{A}^{-1}\mathbf{f}$. In applying the finite element method to our visual surface reconstruction problem, we will obtain large sparse matrices \mathbf{A} . In this appendix we will be concerned with numerical methods which are capable of solving such systems where N is in the range of, say, $10^3 - 10^6$, or larger.

An *iterative method* for solving equations (2) computes a sequence of approximations $\mathbf{u}^{(1)}, \mathbf{u}^{(2)}, \dots$, to the exact solution \mathbf{u} . A new, and hopefully better, approximation is computed at each iteration, but, in general, the exact solution cannot be obtained in a finite number of iterations. If, regardless of the initial approximation $\mathbf{u}^{(0)}$,¹⁴ the *approximation error* (i.e., the difference between the exact solution and the approximations, measured in some appropriate norm) tends to zero as the number of iterations increases, the iterative method is said to be *convergent*, and the rate at which the approximation error tends to zero is called its *rate of convergence*. In order that an iterative method be of practical use, it is important that it be convergent, and that it exhibit a sufficiently large rate of convergence to an approximation of prespecified accuracy. In this appendix, we review a number of the most common iterative methods, and examine their convergence properties and their rates of convergence. References for this material are [Forsythe and Wasow, 1960], [Smith, 1977], [Varga, 1963] and [Young, 1971; Young and Gregory, 1972, 1973].

B.1. Basic Relaxation Methods

Let us assume that \mathbf{A} has nonzero diagonal elements. It will be convenient to express \mathbf{A} as the sum

$$\mathbf{A} = \mathbf{D} - \mathbf{L} - \mathbf{U},$$

¹⁴The trivial initial approximation $\mathbf{u}^{(0)} = \mathbf{0}$ is usually chosen.

where $\mathbf{D} \in \mathbb{R}^{NN}$ is diagonal, $\mathbf{L} \in \mathbb{R}^{NN}$ is strictly lower triangular, and $\mathbf{U} \in \mathbb{R}^{NN}$ is strictly upper triangular. Clearly the equations in (1) can be rewritten in the following form:

$$a_{ii}u_i = - \sum_{\substack{1 \leq j \leq N \\ j \neq i}} a_{ij}u_j + f_i, \quad 1 \leq i \leq N.$$

Next, we will define three basic iterative methods for solving (2), popularly known as *relaxation methods* in the context of the numerical solution of partial differential equations.

Jacobi Relaxation

The *Jacobi method* (or the *method of simultaneous displacements*) is defined by

$$a_{ii}u_i^{(k+1)} = - \sum_{\substack{1 \leq j \leq N \\ j \neq i}} a_{ij}u_j^{(k)} + f_i, \quad 1 \leq i \leq N,$$

which can be written in matrix form as

$$\mathbf{D}\mathbf{u}^{(k+1)} = (\mathbf{L} + \mathbf{U})\mathbf{u}^{(k)} + \mathbf{f},$$

thus giving us the iterative scheme

$$\mathbf{u}^{(k+1)} = \mathbf{D}^{-1}(\mathbf{L} + \mathbf{U})\mathbf{u}^{(k)} + \mathbf{D}^{-1}\mathbf{f}. \quad (3)$$

The matrix

$$\mathbf{G}_J = \mathbf{D}^{-1}(\mathbf{L} + \mathbf{U})$$

is called the *Jacobi iteration matrix* associated with matrix \mathbf{A} .

Clearly, the Jacobi method is a *parallel* method because the elements of the new approximation $\mathbf{u}^{(k+1)}$ may be computed simultaneously and in parallel by a network of processors whose inputs are elements of the old approximation $\mathbf{u}^{(k)}$. As such, it requires the storage of both the old and the new approximations.

Gauss-Seidel Relaxation

Convergence of the Jacobi method is usually very slow. In a closely related method, the so called *Gauss-Seidel method* (or the *method of immediate displacements*), the elements of the new approximation are used in subsequent computations immediately after they become available. This increases the rate of convergence somewhat, but typically less than an order of magnitude. The equations are written as

$$a_{ii}u_i^{(k+1)} = - \sum_{j=1}^{i-1} a_{ij}u_j^{(k+1)} - \sum_{j=i+1}^N a_{ij}u_j^{(k)} + f_i, \quad 1 \leq i \leq N,$$

which, in matrix form, becomes

$$\mathbf{D}\mathbf{u}^{(k+1)} = \mathbf{L}\mathbf{u}^{(k+1)} + \mathbf{U}\mathbf{u}^{(k)} + \mathbf{f}, \quad (4)$$

from which we obtain the iterative scheme defined by

$$\mathbf{u}^{(k+1)} = (\mathbf{D} - \mathbf{L})^{-1} \mathbf{U} \mathbf{u}^{(k)} + (\mathbf{D} - \mathbf{L})^{-1} \mathbf{f}. \quad (5)$$

The *Gauss-Seidel iteration matrix* associated with matrix \mathbf{A} is therefore given by

$$\mathbf{G}_{GS} = (\mathbf{D} - \mathbf{L})^{-1} \mathbf{U}.$$

In the Gauss-Seidel method, we no longer need to store simultaneously the old and new approximations. As they are computed, the new elements can simply displace the old ones. Moreover, since the new values are exploited immediately in subsequent computations we can intuitively expect a higher rate of convergence compared with the Jacobi method. On the other hand, it can easily be seen that the Gauss-Seidel method is inherently a *sequential* method which renders it unsuitable for implementation as a parallel network.

The Successive Overrelaxation (SOR) Method

The convergence of the Gauss-Seidel method may be accelerated by a simple modification. Let us define the *dynamic residual vector*¹⁵ at the k^{th} iteration of the Gauss-Seidel relaxation method as

$$\begin{aligned} \mathbf{r}^{(k)} &= \mathbf{u}^{(k+1)} - \mathbf{u}^{(k)} \\ &= \mathbf{D}^{-1} (\mathbf{L} \mathbf{u}^{(k+1)} + \mathbf{U} \mathbf{u}^{(k)} + \mathbf{f}) - \mathbf{u}^{(k)} \end{aligned}$$

(see (4)). Then, it is evident that the Gauss-Seidel scheme can be written in the form

$$\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \mathbf{r}^{(k)}.$$

If the elements of all successive residual vectors are one-signed (as they usually are when approximating elliptic problems), then it is reasonable to anticipate an acceleration in the convergence of the Gauss-Seidel scheme if the residual vector is scaled by a fixed real number $\omega > 1$ before it is added to $\mathbf{u}^{(k)}$ in each Gauss-Seidel iteration. ω is called the *relaxation parameter*. This idea leads to one of the most successful iterative methods for solving systems of linear equations currently available, the *successive overrelaxation method*. The method may be defined by

$$\begin{aligned} \mathbf{u}^{(k+1)} &= \mathbf{u}^{(k)} + \omega \mathbf{r}^{(k)} \\ &= \mathbf{u}^{(k)} + \omega [\mathbf{D}^{-1} (\mathbf{L} \mathbf{u}^{(k+1)} + \mathbf{U} \mathbf{u}^{(k)} + \mathbf{f}) - \mathbf{u}^{(k)}], \end{aligned}$$

which we can manipulate into the form

$$\mathbf{u}^{(k+1)} = (\mathbf{I} - \omega \mathbf{D}^{-1} \mathbf{L})^{-1} [(1 - \omega) \mathbf{I} + \omega \mathbf{D}^{-1} \mathbf{U}] \mathbf{u}^{(k)} + (\mathbf{I} - \omega \mathbf{D}^{-1} \mathbf{L})^{-1} \omega \mathbf{D}^{-1} \mathbf{f}. \quad (6)$$

The *successive overrelaxation iteration matrix* is therefore given by

$$\mathbf{G}_\omega = (\mathbf{I} - \omega \mathbf{D}^{-1} \mathbf{L})^{-1} [(1 - \omega) \mathbf{I} + \omega \mathbf{D}^{-1} \mathbf{U}].$$

¹⁵The residual vector is also called the *correction* or *displacement vector*. Note that the *residual* of an equation is the amount by which the equation fails to be satisfied.

Three cases arise:

- (i) if $\omega > 1$ the scheme is called overrelaxation,¹⁶
- (ii) if $\omega = 1$ we obtain the Gauss-Seidel method, and
- (iii) if $0 < \omega < 1$ the scheme is termed underrelaxation.

Conditions for Convergence

The Jacobi, Gauss-Seidel, and successive overrelaxation methods are particular instances of the general *stationary iterative method*

$$\mathbf{u}^{(k+1)} = \mathbf{G}\mathbf{u}^{(k)} + \mathbf{c}, \quad (7)$$

where the iteration matrix \mathbf{G} is taken to be \mathbf{G}_J , \mathbf{G}_{GS} , and \mathbf{G}_w respectively, and \mathbf{c} is a known column vector.¹⁷ By subtracting $\mathbf{u} = \mathbf{G}\mathbf{u} + \mathbf{c}$ from (7), we can obtain an expression for the errors $\mathbf{e}^{(k)} = \mathbf{u}^{(k)} - \mathbf{u}$ of successive approximations,

$$\begin{aligned} \mathbf{e}^{(k+1)} &= \mathbf{G}\mathbf{e}^{(k)} \\ &= \mathbf{G}^{k+1}\mathbf{e}^{(0)}. \end{aligned} \quad (8)$$

The sequence of approximations converges to the solution \mathbf{u} if $\lim_{k \rightarrow \infty} \mathbf{e}^{(k)} = \mathbf{0}$, which will clearly be the case if and only if $\lim_{k \rightarrow \infty} \mathbf{G}^k = \mathbf{0}$, since $\mathbf{u}^{(0)}$ (and hence $\mathbf{e}^{(0)}$) is arbitrary. Let \mathbf{G} have eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_N$, and assume that the corresponding eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N$ are linearly independent. Then the initial error vector can be expressed uniquely as the linear combination

$$\mathbf{e}^{(0)} = \sum_{i=1}^N c_i \mathbf{v}_i.$$

But, by (8),

$$\begin{aligned} \mathbf{e}^{(k)} &= \mathbf{G}^k \mathbf{e}^{(0)} \\ &= \sum_{i=1}^N c_i \mathbf{G}^k \mathbf{v}_i \\ &= \sum_{i=1}^N c_i \lambda_i^k \mathbf{v}_i. \end{aligned}$$

It follows that in the limit, $\mathbf{e}^{(k)}$ will be zero for an arbitrary initial approximation if and only if $|\lambda_i| < 1$, for $1 \leq i \leq N$. Thus, we have the following theorem.

¹⁶It should be noted that there exist classes of matrices (which arise from many first and second order partial differential equations) for which the optimal value of ω , yielding the largest rate of convergence, may be determined analytically (see e.g. [Young, 1972], [Young and Gregory, 1973]). Often, the convergence may be adequately accelerated by not necessarily optimal values of ω chosen empirically. Of course, it is possible to vary ω from iteration to iteration or from one equation to the next. A number of these modified methods have been studied in the literature (see e.g. [Varga, 1962] or [Young, 1971; Young and Gregory, 1972, 1973]).

¹⁷The method is obtained by writing the original system $\mathbf{A}\mathbf{u} = \mathbf{f}$ as $\mathbf{u} = \mathbf{G}\mathbf{u} + \mathbf{c}$, and is referred to as being stationary because \mathbf{G} is fixed for all iterations.

Theorem 1 — necessary and sufficient condition for convergence of the stationary iterative method

The stationary iterative method (7) is convergent if and only if

$$\rho(\mathbf{G}) = \max_{i=1}^N |\lambda_i(\mathbf{G})| < 1,$$

where $\rho(\mathbf{G})$ is called the spectral radius of \mathbf{G} .¹⁸

Theorem 1 is mainly of theoretical value because, in practice, it is difficult to determine the eigenvalues of \mathbf{G} . Fortunately, we have a useful corollary giving sufficient conditions for convergence. The corollary results from the observation that for some matrix norm $\|\mathbf{G}\|$ which is consistent¹⁹ with a vector norm $\|\mathbf{v}_i\|$,

$$\|\mathbf{e}^{(k)}\| \leq \|\mathbf{G}^k\| \|\mathbf{e}^{(0)}\| \leq \|\mathbf{G}\|^k \|\mathbf{e}^{(0)}\|. \quad (9)$$

Hence we obtain the following corollary to Theorem 1.

Corollary 1 — sufficient condition for convergence of the stationary iterative method

*If $\|\mathbf{G}\| < 1$, then the stationary iterative method (7) is convergent.*²⁰

Suppose that the stationary iterative method is convergent. It can be shown (see [Varga, 1962] or [Young, 1971]) that

$$\lim_{k \rightarrow \infty} \|\mathbf{G}^k\|_{L_2}^{1/k} = \rho(\mathbf{G}).$$

Hence, from (9) we have that, for large k ,

$$\|\mathbf{e}^{(k)}\|_{L_2} \approx \rho(\mathbf{G})^k \|\mathbf{e}^{(0)}\|_{L_2}.$$

Thus, in a certain sense, $\rho(\mathbf{G})$ is a measure of the rate of convergence of the iterative method and therefore, like convergence itself, depends on the eigenvalues of \mathbf{G} .

We illustrate an application of Corollary 1 in obtaining a simple but important sufficient condition for the convergence of Jacobi relaxation. The Jacobi iteration matrix \mathbf{G}_J consists of the elements

$$g_{ij} = \frac{a_{ij}}{a_{ii}}, \quad i \neq j; \\ g_{ii} = 0.$$

Therefore the L_∞ norm of \mathbf{G}_J is given by

$$\|\mathbf{G}_J\|_\infty = \max_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N \frac{|a_{ij}|}{|a_{ii}|}.$$

¹⁸It can be shown (see e.g. [Varga, 1962]) that the theorem holds without the independence assumption on the eigenvectors.

¹⁹If a matrix norm and a vector norm satisfy the relation $\|\mathbf{A}\mathbf{u}\| \leq \|\mathbf{A}\| \|\mathbf{u}\|$ for any \mathbf{A} and \mathbf{u} , then the two norms are said to be *consistent* [Dahlquist and Björck, 1974, pg. 175].

²⁰The condition is not a necessary one because we can have the case that $\|\mathbf{G}\| > 1$ when $\rho(\mathbf{G}) < 1$.

Clearly, if $|a_{ii}| > \sum_{\substack{1 \leq j \leq N \\ j \neq i}} |a_{ij}|$, $\|G_J\|_\infty < 1$, and Jacobi relaxation will converge by corollary 1. A much more general result may be obtained. We begin by defining two important properties that A may possess.

Definition 1 — (weakly) diagonally dominant matrix

A matrix A of order N is said to be diagonally dominant if

$$|a_{ii}| > \sum_{\substack{1 \leq j \leq N \\ j \neq i}} |a_{ij}|, \quad 1 \leq i \leq N.$$

The matrix is said to be weakly diagonally dominant if the $>$ relation in the above inequality can be replaced by \geq in some, but not all, of the equations.

Definition 2 — irreducible matrix

A matrix A of order N is irreducible if and only if $N = 1$, or if $N > 1$ and for any i and j such that $1 \leq i, j \leq N$ and $i \neq j$ either $a_{ij} = 0$, or there exist i_1, i_2, \dots, i_k such that $a_{i,i_1} a_{i_1,i_2} \dots a_{i_k,j} \neq 0$.²¹

It can be shown that if A is irreducible and has weak diagonal dominance, then it is nonsingular, and if in addition it is symmetric and has non-negative diagonal elements, then it is positive definite. The general theorem is given next (For a proof see [Varga, 1962, pg. 73]).

Theorem 2 — sufficient conditions for convergence of the Jacobi and Gauss-Seidel relaxation

Let A be either a diagonally dominant or an irreducible and weakly diagonally dominant matrix. Then, both the associated Jacobi and Gauss-Seidel relaxation methods of (4) and (5) are convergent.

Next, we turn our attention to the successive overrelaxation method. Since the inverse of a triangular matrix is also a triangular matrix, and its determinant is equal to the product of its diagonal elements, we have

$$\det(G_\omega) = \det[(I - \omega D^{-1}L)^{-1}] \det[(1 - \omega)I + \omega D^{-1}U] = (1 - \omega)^N.$$

Since $\det(G_\omega) = \prod_{i=1}^N \lambda_i$, it follows that

$$\max_{i=1}^N |\lambda_i| \geq |1 - \omega|,$$

which, by theorem 1, leads to the following.

Theorem 3 — convergence of the SOR method

$$\rho(G_\omega) \geq |\omega - 1|$$

²¹The term *irreducible matrix* was introduced by Frobenius for matrices which (informally speaking) correspond to systems of equations whose solutions cannot be reduced to the solution of two systems of lower order. One generally obtains irreducible matrices when discretizing boundary value problems over connected domains.

therefore the successive overrelaxation method (6) can converge only if $0 < \omega < 2$.

A set of necessary and sufficient conditions for convergence for the successive overrelaxation method are stated in the following theorem.

Theorem 4 — convergence of SOR method for symmetric, positive definite A

If A is real, symmetric matrix with positive diagonal elements, then the successive overrelaxation method (5) is convergent if and only if $0 < \omega < 2$ and A is positive definite.

The same conditions for convergence hold for the Gauss-Seidel method since, by definition, it is a special case of successive overrelaxation.

Corollary 2 — convergence of Gauss-Seidel relaxation for symmetric, positive definite A

If A is a symmetric matrix with positive diagonal elements, then the Gauss-Seidel method is convergent if and only if A is positive definite.

It is important to note that the same statement cannot be made about the Jacobi method.

B.2. Basic Gradient Methods

In this section we investigate a class of iterative methods which are naturally associated with optimization theory. These are the so called *gradient methods* which, in their full generality, are iterative techniques for minimizing nonlinear functionals. They may also be thought of as methods for solving systems of linear equations for the special case where the functional to be minimized is a quadratic form.

Assume that $A \in \mathbb{R}^{N \times N}$ is symmetric. Now suppose that we attempt to solve the following *unconstrained minimization problem* involving the quadratic form $\mathcal{E}(\mathbf{v})$

$$\mathcal{E}(\mathbf{u}) = \inf_{\mathbf{v} \in \mathbb{R}^N} \mathcal{E}(\mathbf{v}) = \frac{1}{2}(\mathbf{v}, A\mathbf{v}) - (\mathbf{f}, \mathbf{v}),$$

where $(\cdot, \cdot): \mathbb{R}^N \times \mathbb{R}^N \mapsto \mathbb{R}$ denotes the familiar Euclidean inner product. From the well-known theorem of optimization theory (see, e.g. [Luenberger, 1973]), the *gradient vector* of $\mathcal{E}(\mathbf{u})$

$$\nabla \mathcal{E}(\mathbf{u}) = A\mathbf{u} - \mathbf{f},$$

vanishes at a minimum \mathbf{u} and, moreover, the minimum exists and is unique if the *Hessian matrix*

$$\left[\frac{\partial^2 \mathcal{E}(\mathbf{u})}{\partial u_i \partial u_j} \right] = A.$$

is positive definite.

Thus, for symmetric, positive definite A , solving the minimization problem is equivalent to solving the system of linear equations $A\mathbf{u} = \mathbf{f}$ and, consequently, relaxation methods can be thought of as being methods for descending to the minimum \mathbf{u} of a quadratic functional.

Gradient Descent

Consider the iterative method $\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \alpha^{(k)}\mathbf{d}^{(k)}$, where at each iteration, we take a step in the direction of the vector $\mathbf{d}^{(k)}$. To minimize $\mathcal{E}(\mathbf{u})$ quickly, we should move in the *direction of steepest descent*, which is given by the negative gradient; that is, $\mathbf{d}^{(k)} = -\nabla \mathcal{E}(\mathbf{u}^{(k)}) = \mathbf{f} - \mathbf{A}\mathbf{u}^{(k)} = \mathbf{r}^{(k)}$, where $\mathbf{r}^{(k)}$ is the familiar residual vector. Thus we obtain the nonstationary iterative method defined by

$$\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \alpha^{(k)}\mathbf{r}^{(k)}. \quad (10)$$

It seems reasonable to choose the step size at each iteration $\alpha^{(k)}$ so as to minimize $\mathcal{E}(\mathbf{u}^{(k)} + \alpha^{(k)}\mathbf{r}^{(k)})$. The appropriate value can easily be shown to be

$$\alpha^{(k)} = \frac{(\mathbf{r}^{(k)}, \mathbf{r}^{(k)})}{(\mathbf{r}^{(k)}, \mathbf{A}\mathbf{r}^{(k)})}. \quad (11)$$

On the other hand, if we fix $\alpha^{(k)} = \alpha$ for all iterations, then (10) becomes

$$\mathbf{u}^{(k+1)} = (\mathbf{I} - \alpha\mathbf{A})\mathbf{u}^{(k)} + \alpha\mathbf{f}, \quad (12)$$

which can be identified as a stationary iterative method with iteration matrix $\mathbf{G} = \mathbf{I} - \alpha\mathbf{A}$.

The Conjugate Gradient Method

Hestenes and Stiefel [1952] introduced the *conjugate gradient method*, a modification of the method of gradient descent. The method is based on determining vectors $\mathbf{d}^{(0)}, \mathbf{d}^{(1)}, \dots, \mathbf{d}^{(n-1)}$ which are pairwise conjugate in the sense that $(\mathbf{d}^{(i)}, \mathbf{A}\mathbf{d}^{(j)}) = 0$ for $i \neq j$. The ease in applying the method derives from the fact that these vectors may also be determined iteratively. With the residual vector at the k^{th} iteration given by $\mathbf{r}^{(k)} = \mathbf{f} - \mathbf{A}\mathbf{u}^{(k)}$ and with $\mathbf{d}^{(0)} = \mathbf{r}^{(0)}$, the algorithm for determining the $\mathbf{d}^{(k)}$ and $\mathbf{u}^{(k)}$ is as follows.

$$\begin{cases} \mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \frac{(\mathbf{r}^{(k)}, \mathbf{d}^{(k)})}{(\mathbf{d}^{(k)}, \mathbf{A}\mathbf{d}^{(k)})} \mathbf{d}^{(k)}, & 0 \leq k \leq N-1; \\ \mathbf{d}^{(k)} = \mathbf{r}^{(k)} - \frac{(\mathbf{r}^{(k)}, \mathbf{A}\mathbf{d}^{(k-1)})}{(\mathbf{d}^{(k-1)}, \mathbf{A}\mathbf{d}^{(k-1)})} \mathbf{d}^{(k-1)}, & 1 \leq k \leq N-1. \end{cases}$$

Conjugacy of the vectors can be verified along with the fact that $(\mathbf{r}^{(i)}, \mathbf{r}^{(j)}) = 0$ for $i \neq j$. This implies that $\mathbf{r}^{(k)} = 0$ for some $k \leq N$. Therefore, in the absence of roundoff errors, the method converges in at most N iterations. Of course, this property is of no real value to us because we must deal with cases where N is very large. Nevertheless, for N large, typically $\mathbf{u}^{(k)} \approx \mathbf{u}$ for $k \ll N$ and the algorithm may be used in the iterative spirit. The conjugate gradient method is certainly the most expensive of the algorithms discussed both in terms of space (since the vectors $\mathbf{u}^{(k)}$, $\mathbf{d}^{(k)}$, $\mathbf{r}^{(k)}$, and $\mathbf{A}\mathbf{d}^{(k)}$ must be stored) and in terms of the number of operations to complete one iteration. While it is true that, for model problems, the number of iterations required to reduce the error by a specified amount is usually considerably less than for the other methods, the conjugate gradient method seems to exceed at least the successive overrelaxation method (with optimal ω) in total number of operations required [see e.g. Young and Gregory, 1973, pg. 1071].

Convergence and Comparisons to Relaxation

It is easy to show [Luenberger, 1973] that gradient descent (10) with (11) is convergent for a positive definite matrix A to the solution $u = A^{-1}f$ where the quadratic form \mathcal{E} is minimized. On the other hand, we must be a little more careful with the fixed- α descent algorithm (12). The eigenvalues λ_i of its iteration matrix G are related to the eigenvalues λ'_i of A , all of which are positive (since A is positive definite), by $\lambda_i = 1 - \alpha\lambda'_i$. Therefore, according to theorem 1, we obtain the following.

Theorem 5 — necessary and sufficient condition for convergence

For a positive definite matrix A , the fixed- α method of equation (12) is convergent if and only if $0 < \alpha < \frac{2}{\rho(A)}$.

Of course, convergence is quickest for that α which minimizes $\rho(G)$, which often cannot be determined in practice.

By comparing (12) with the Jacobi method (3), we can convince ourselves that the two methods become identical when A is positive definite, has identical elements on its main diagonal (i.e. $D = aI$), and $\alpha = 1/a$. Moreover, Forsythe and Wasow [1960, pg. 239] (see also [Milne, 1970]) show that the Gauss-Seidel and successive overrelaxation methods are also subject to interpretations as descent methods. In these cases, however, we have not one, but a set of direction vectors which turn out to be \mathbf{x}_i , the coordinate unit vectors of \mathbb{R}^N . During each iteration, we take a sequence of steps of different sizes in each of these directions. The step sizes are such that $\mathcal{E}(\mathbf{u}^{(k)} + \alpha_i^{(k)}\mathbf{x}_i)$ is minimized, and are given by $\alpha_i^{(k)} = r_i^{(k)}/a_{ii}$ for the Gauss-Seidel method and $\alpha_i^{(k)} = \omega r_i^{(k)}/a_{ii}$ for the successive overrelaxation method, where $r_i^{(k)}$ is the i^{th} element of the residual vector at iteration k .

The computation of an optimal $\alpha^{(k)}$ at each iteration according to (11) requires that $\mathbf{r}^{(k)}$ be stored and that $A\mathbf{r}^{(k)}$ be evaluated. This doubles the amount of work required per iteration in comparison with a fixed- α algorithm or Jacobi relaxation. This raises the question: which is better in the long run; N iterations of gradient descent, or $2N$ iterations of the fixed- α or Jacobi relaxation? To quantitatively decide the issue, a convergence analysis ought to be attempted. This is problem-dependent and is generally difficult to do, so we will simply note that Forsythe and Wasow [1960, pg. 225] do not recommend the optimization of α and, moreover, refer to a result by Stiefel [1955] indicating that it is, at best, a short-sighted strategy. Interestingly enough, Grimson [1981a] used the optimal- α algorithm in his implementation of the surface interpolation algorithm (with a minor modification to make certain that the fixed constraints are never modified, thus, in effect, treating them as essential boundary conditions). Considering the statements of Forsythe and Wasow, it is not surprising that extremely slow convergence was observed in spite of the extra work expended at each iteration to compute the optimal α .

C. LOCAL FOURIER ANALYSIS OF RELAXATION

In this appendix, we present the details of a local Fourier analysis of the Gauss-Seidel relaxation and obtain the smoothing factor for this method. The analysis involves studying separately the convergence of the high-frequency Fourier components. Since these components have short coupling ranges, we can perform the analysis in the interior of $\bar{\Omega}$, ignoring the effects of the boundary and the constraints.

According to equation (4.9), the minimizing displacement $u_{i,j}$ at an interior node (i,j) is related to the other nodes by

$$\begin{aligned} 20u_{i,j} - 8(u_{i-1,j} + u_{i+1,j} + u_{i,j-1} + u_{i,j+1}) \\ + 2(u_{i-1,j-1} + u_{i+1,j-1} + u_{i-1,j+1} + u_{i+1,j+1}) \\ + 1(u_{i-2,j} + u_{i+2,j} + u_{i,j-2} + u_{i,j+2}) = 0 \end{aligned} \quad (1)$$

(for convenience, we have suppressed the superscripts h). According to the discussion in Appendix B, at iteration k of the Gauss-Seidel relaxation method, $v_{i,j}^{(k)}$ is replaced by a new value $v_{i,j}^{(k+1)}$ such that

$$\begin{aligned} 20v_{i,j}^{(k+1)} - 8(v_{i-1,j}^{(k+1)} + v_{i+1,j}^{(k)} + v_{i,j-1}^{(k+1)} + v_{i,j+1}^{(k)}) \\ + 2(v_{i-1,j-1}^{(k+1)} + v_{i+1,j-1}^{(k+1)} + v_{i-1,j+1}^{(k)} + v_{i+1,j+1}^{(k)}) \\ + 1(v_{i-2,j}^{(k+1)} + v_{i+2,j}^{(k)} + v_{i,j-2}^{(k+1)} + v_{i,j+2}^{(k)}) = 0. \end{aligned} \quad (2)$$

The errors of the approximation at iteration k and iteration $k+1$ are given by

$$e_{i,j}^{(k)} = u_{i,j} - v_{i,j}^{(k)} \quad \text{and} \quad e_{i,j}^{(k+1)} = u_{i,j} - v_{i,j}^{(k+1)}$$

respectively. Subtracting (2) from (1), we obtain

$$\begin{aligned} 20e_{i,j}^{(k+1)} - 8(e_{i-1,j}^{(k+1)} + e_{i+1,j}^{(k)} + e_{i,j-1}^{(k+1)} + e_{i,j+1}^{(k)}) \\ + 2(e_{i-1,j-1}^{(k+1)} + e_{i+1,j-1}^{(k+1)} + e_{i-1,j+1}^{(k)} + e_{i+1,j+1}^{(k)}) \\ + 1(e_{i-2,j}^{(k+1)} + e_{i+2,j}^{(k)} + e_{i,j-2}^{(k+1)} + e_{i,j+2}^{(k)}) = 0, \end{aligned} \quad (3)$$

Suppose that the error consists of only a single spatial Fourier component $\bar{\omega} = [\omega_1, \omega_2]$. Then the errors at node (i,j) before and after the k^{th} iteration are given by

$$e_{i,j}^{(k)} = A_{\bar{\omega}}^{(k)} e^{\iota(\omega_1 i + \omega_2 j)} \quad \text{and} \quad e_{i,j}^{(k+1)} = A_{\bar{\omega}}^{(k+1)} e^{\iota(\omega_1 i + \omega_2 j)} \quad (4)$$

respectively, where $\iota = \sqrt{-1}$.

Substituting (4) into (3), dividing through by $e^{\iota(\omega_1 i + \omega_2 j)}$, and collecting terms pertaining to the same iteration, we obtain

$$A_{\vec{\omega}}^{(k)} \left(-8(e^{i\omega_1} + e^{i\omega_2}) + 2(e^{i(\omega_1+\omega_2)} + e^{i(-\omega_1+\omega_2)}) + (e^{2i\omega_1} + e^{2i\omega_2}) \right) \\ + A_{\vec{\omega}}^{(k+1)} \left(20 - 8(e^{-i\omega_1} + e^{-i\omega_2}) + 2(e^{i(\omega_1-\omega_2)} + e^{i(-\omega_1-\omega_2)}) + (e^{-2i\omega_1} + e^{-2i\omega_2}) \right) = 0.$$

The amplification of the $\vec{\omega}$ component is then given by

$$\mu_{GS}(\vec{\omega}) = \left| \frac{A_{\vec{\omega}}^{(k+1)}}{A_{\vec{\omega}}^{(k)}} \right| \\ = \left| \frac{8(e^{i\omega_1} + e^{i\omega_2}) - 2(e^{i(\omega_1+\omega_2)} + e^{i(-\omega_1+\omega_2)}) - (e^{2i\omega_1} + e^{2i\omega_2})}{20 - 8(e^{-i\omega_1} + e^{-i\omega_2}) + 2(e^{i(\omega_1-\omega_2)} + e^{i(-\omega_1-\omega_2)}) + (e^{-2i\omega_1} + e^{-2i\omega_2})} \right|.$$

Let $|\vec{\omega}| = \max(|\omega_1|, |\omega_2|)$. The Gauss-Seidel smoothing factor is defined as the smallest amplification attained for a high-frequency component of the error; that is, a component which is visible on the fine grid, but is aliased on the coarse grid:

$$\bar{\mu}_{GS} = \max_{\frac{\pi}{2} \leq |\vec{\omega}| \leq \pi} \mu_{GS}(\vec{\omega}).$$

Evaluating this expression numerically, we obtain $\bar{\mu}_{GS} \approx 0.8$ (for $\omega_1 = 1.6$ and $\omega_2 = 0.3$).

REFERENCES

- Abramowitz, M., and Stegun, I.A. (eds.), [1965]
Handbook of Mathematical Functions, Dover, New York.
- Adams, R.A., [1975]
Sobolev Spaces, Academic Press, New York.
- Agmon, S., [1965]
Lectures on Elliptic Boundary Value Problems, Van Nostrand, New Jersey.
- Ahlberg, J.H., Nilson, E.N., and Walsh, J.L., [1967]
The Theory of Splines and their Applications, New York.
- Babuska, I., [1973]
 "The finite element method with penalty functions," *Math. Comp.*, **27**, 221-228.
- Barlow, H.B., [1979]
 "Reconstructing the visual image in space and time," *Nature, Lond.*, **279**, 189-190.
- Brady, J.M., and Horn, B.K.P., [1981]
 Rotationally symmetric operators for surface interpolation, MIT AI Laboratory, Cambridge, MA, AI Memo No. 654.
- Brandt, A., [1977a]
 "Multi-level adaptive solutions to boundary value problems," *Math. Comp.*, **31**, 333-390.
- Brandt, A., [1977b]
 "Multi-level adaptive techniques (MLAT) for partial differential equations: ideas and software," *Mathematical Software III*, Rice, J.R. (ed.), Academic Press, New York, 277-318.
- Brandt, A., and Dinar, N., [1979]
 "Multi-grid solutions to elliptic flow problems," *Numerical Methods for Partial Differential Equations*, Parter, S.V. (ed.), Academic Press, New York, 53-147.
- Brooks, R.A., [1981]
 "Symbolic reasoning among 3-D models and 2-D images," *Artificial Intelligence*, **17**, 285-348.
- Campbell, F.W., and Robson, J.G., [1968]
 "Application of Fourier analysis to the visibility of gratings," *Jour. Physiol. Lond.*, **197**, 551-566.
- Ciarlet, P.G., [1978]
The Finite Element Method for Elliptic Problems, North Holland, Amsterdam.
- Collatz, L., [1966]
The Numerical Treatment of Differential Equations; Third edition, Springer-Verlag, New York.
- Courant, R., [1943]
 "Variational methods for the solution of problems of equilibrium and vibrations," *Bull. Amer. Math. Soc.*, **49**, 1-23.
- Courant, R., and Hilbert, D., [1953]
Methods of Mathematical Physics, Vol. I, Interscience, London.
- Crick, F.H.C., Marr, D., and Poggio, T., [1981]
 "An information processing approach to understanding the visual cortex," *The Organization of the Cerebral Cortex*, Schmitt, F.O., (ed.), MIT Press, Cambridge, MA, 505-533.
- Dahlquist, G., and Björck, A. (Anderson, N., trans.), [1974]
Numerical Methods, Prentice-Hall, Englewood Cliffs, NJ.

- Davis, L.S., and Rosenfeld, A., [1981]
 "Cooperating processes for low-level vision: a survey," *Artificial Intelligence*, **17**, 245-263.
- Duchon, J., [1976]
 "Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces," *R.A.I.R.O. Analyse Numérique*, **10**, 5-12.
- Duchon, J., [1977]
 "Splines minimizing rotation-invariant semi-norms in Sobolev spaces," *Constructive Theory of Functions of Several Variables*, Dodd, A., and Eckmann, B., (eds.), Springer-Verlag, Berlin, 85-100.
- Fahle, M., and Poggio, T., [1981]
 "Visual hyperacuity: spatiotemporal interpolation in human vision," *Proc. Roy. Soc. Lond., B*, **213**, 451-477.
- Faugeras, O., and Berthod, M., [1981]
 "Improving consistency and reducing ambiguity in stochastic labeling: an optimization approach," *IEEE Trans. Pat. Anal. Machine Intel.*, **PAMI-3**, 412-424.
- Finlayson, B.A., [1972]
The Method of Weighted Residuals and Variational Principles, Academic Press, New York.
- Forsythe, G.E., and Wasow, W.R., [1960]
Finite Difference Methods for Partial Differential Equations, Wiley, New York.
- Gladwell, I., and Wait, R., (eds.), [1979]
A Survey of Numerical Methods for Partial Differential Equations, Clarendon, Oxford.
- Grimson, W.E.L., [1981a]
From Images to Surfaces: A Computational Study of the Human Early Visual System, MIT Press, Cambridge, MA.
- Grimson, W.E.L., [1981b]
 "A computer implementation of a theory of human stereo vision," *Phil. Trans. Roy. Soc. Lond., B*, **292**, 217-253.
- Grimson, W.E.L., [1982a]
 "A computational theory of visual surface interpolation," *Phil. Trans. Roy. Soc. Lond., B*, in press.
- Grimson, W.E.L., [1982b]
 The implicit constraints of the primal sketch, MIT AI Laboratory, Cambridge, MA, AI Memo No. 663, to appear.
- Harder, R.L., and Desmarais, R.N., [1972]
 "Interpolation using surface splines," *Jour. Aircraft*, **9**, 189-191.
- Hestenes, M.R., and Stiefel, E., [1952]
 "Method of conjugate gradients for solving linear systems," *Jour. Res. National Bureau Standards*, **49**, 409-436.
- Hilbert, D., and Cohn-Vossen, S., [1952]
Geometry and the Imagination, Chelsea, New York.
- Hildreth, E.C., [1980]
 Implementation of a theory of edge detection, MIT AI Laboratory, Cambridge, MA, AI-TR-579.
- Horn, B.K.P., [1974]
 "Determining lightness from an image," *Computer Graphics and Image Processing*, **3**, 111-299.
- Horn, B.K.P., [1975]
 "Obtaining shape from shading information," *The Psychology of Computer Vision*, Winston, P.H., (ed.), McGraw-Hill, New York, 115-155.

- Horn, B.K.P., and Schunck, B.G., [1981]
 "Determining optical flow," *Artificial Intelligence*, 17, 185-203.
- Hummel, R., and Zucker, S.W., [1980]
 On the foundations of relaxation labeling, Computer Vision and Graphics Laboratory, McGill University, Montreal, TR-80-7, to appear in *IEEE Trans. Pat. Anal. Machine Intel.*
- Ikeuchi, K., [1981]
 "Determining surface orientations of specular surfaces by using the photometric stereo method," *IEEE Trans. Pat. Anal. Machine Intel.*, 661-669.
- Ikeuchi, K., and Horn, B.K.P., [1981]
 "Numerical shape from shading and occluding boundaries," *Artificial Intelligence*, 17, 141-184.
- Johansson, G., [1975]
 "Visual motion perception," *Scientific American*, 232, 76-88.
- Julesz, B., [1971]
Foundations of Cyclopean Perception, University of Chicago Press, Chicago, IL.
- Kanade, T.,
 "Recovery of the three-dimensional shape of an object from a single view," *Artificial Intelligence*, 17, 409-460.
- Kender, J.R., [1980]
 Shape from Texture, Ph.D. Thesis, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA.
- Landau, L., and Lifshitz, E., [1970]
Theory of Elasticity, Pergamon Press, Oxford.
- Luenberger, D.G., [1973]
Introduction to Linear and Nonlinear Programming, Addison-Wesley, Reading, MA.
- Mansfield, E.H., [1964]
The Bending and Stretching of Plates, Macmillan, New York.
- Mayhew, J.E.W., and Frisby, J.P., [1980]
 "The computation of binocular edges," *Perception*, 9, 69-86.
- Mayhew, J.E.W., and Frisby, J.P., [1981]
 "Psychophysical and computational studies toward a theory of human stereopsis," *Artificial Intelligence*, 17, 349-385.
- Marr, D., [1976]
 "Early processing of visual information," *Phil. Trans. Roy. Soc. Lond., B*, 275, 483-534.
- Marr, D., [1977]
 "Analysis of occluding contour," *Proc. R. Soc. Lond. B*, 197, 441-475.
- Marr, D., [1982]
Vision: A Computational Investigation into the Human Representation and Processing of Visual Information, Freeman, San Francisco, CA.
- Marr, D., and Hildreth, E.C., [1980]
 "Theory of edge detection," *Proc. R. Soc. Lond. B*, 207, 187-217.
- Marr, D., and Nishihara, H.K., [1978]
 "Representation and recognition of the spatial organization of three-dimensional shapes," *Proc. Roy. Soc. Lond., B*, 200, 269-294.
- Marr, D., and Poggio, T., [1977]
 "From understanding computation to understanding neural circuitry," *Neuronal Mechanisms in Visual Perception. Neurosciences Research Program Bulletin; Poppe, E. et al., (eds).*, Vol. 15, 470-488.

- Marr, D., and Poggio, T., [1979]
 "A theory of human stereo vision," *Proc. R. Soc. Lond. B*, 204, 301-328.
- Marr, D., Poggio, T., and Hildreth, E.C., [1980]
 "The smallest channel in early human vision," *Jour. Opt. Soc. Amer.*, 70, 868-870.
- Meinguet, J., [1979a]
 "An intrinsic approach to multivariate spline interpolation at arbitrary points," *Polynomial and Spline Approximation: Theory and Applications*, Sahney, B.N. (ed.), Reidel, Dordrecht, Holland, 163-190.
- Meinguet, J., [1979b]
 "Multivariate interpolation at arbitrary points made simple," *Jour. Applied Math. and Physics (ZAMP)*, 30, 292-304.
- Mikhlin, S.G., [1964]
Variational Methods in Mathematical Physics, Pergamon Press, Elmsford.
- Milne, W.E., [1970]
Numerical Solution of Differential Equations; Second Edition, Dover, New York.
- Nevatia, R., and Binford, T.O., [1977]
 "Description and recognition of curved objects," *Artificial Intelligence*, 8, 77-98.
- Nishihara, H.K., [1981]
 "Intensity, visible surface, and volumetric representations," *Artificial Intelligence*, 17, 265-284.
- Oden, J.T., [1979]
Applied Functional Analysis, Prentice-Hall, Englewood Cliffs, NJ.
- Oden, J.T., and Reddy, J.N., [1976]
An Introduction to the Mathematical Theory of the Finite Element Method, Wiley, New York.
- Ortega, J.M., and Rheinboldt, W.C., [1970]
Iterative Solution of Nonlinear Equations in Several Variables, Academic Press, New York.
- Rashbass, C., and Westheimer, G., [1961]
 "Disjunctive eye movements," *Jour. Physiol. Lond.*, 159, 339-360.
- Rektorys, K., (ed.), [1969]
Survey of Applicable Mathematics, Iliffe, London.
- Rektorys, K., [1980]
Variational Methods in Mathematics, Science, and Engineering; Second edition, Reidel, Dordrecht, Holland.
- Riggs, L.A., and Niehl, E.W., [1960]
 "Eye movements recorded during convergence and divergence," *Jour. Opt. Soc. Amer.*, 50, 913-920.
- Ritz, W., [1908]
 "Ueber eine neue Methode zur Loesung gewisser Variationsprobleme der mathematischen Physik," *Jour. Reine Angew. Math.*, 135.
- Rosenfeld, A., Hummel, R., and Zucker, S.W., [1976]
 "Scene labeling by relaxation operations," *IEEE Trans. Systems, Man, Cybernetics*, 6, 420-433.
- Schwartz, L., [1966]
Théorie des distributions, Hermann, Paris.
- Smith, G.D., [1977]
Numerical Solution of Partial Differential Equations; second edition, Clarendon Press, Oxford.
- Southwell, R.V., [1946]
Relaxation Methods in Theoretical Physics, Oxford University Press, Oxford.

- Stevens, K.A., [1981]
 "The visual interpretation of surface contours," *Artificial Intelligence*, 17, 47-73.
- Stiefel, E., [1955]
 "Relaxationsmethoden bester Strategie zur Lösung linearer Gleichungssysteme," *Comment. Math. Helv.*, 29, 157-179.
- Straug, G., and Fix, G.J., [1973]
An Analysis of the Finite Element Method, Prentice-Hall, Englewoods Cliffs, NJ.
- Ullman, S., [1979a]
The Interpretation of Visual Motion, MIT Press, Cambridge, MA.
- Ullman, S., [1979b]
 "Relaxation and constrained optimization by local processes," *Computer Graphics and Image Processing*, 10, 115-125.
- Varga, I.S., [1962]
Matrix Iterative Analysis, Prentice-Hall, Englewoods Cliffs, NJ.
- Waltz, D., [1975]
 "Understanding line drawings of scenes with shadows," *The Psychology of Computer Vision*, Winston, P.H. (ed.), McGraw-Hill, New York, 19-91.
- Wallach, H., [1959]
 "The perception of motion," *Scientific American*, 210, 56-60.
- Wallach, H., and O'Connell, D.N., [1953]
 "The kinetic depth effect," *J. Exp. Psych.*, 52, 571-578.
- Westheimer, G., [1976]
 "Diffraction theory and visual hyperacuity," *Amer. Jour. Optom. Physiol. Opt.*, 53, 362-364.
- Westheimer, G., and McKee, S.P., [1975]
 "Visual acuity in the presence of retinal-image motion," *Jour. Opt. Soc. Amer.*, 65, 847-850.
- Westheimer, G., and McKee, S.P., [1977]
 "Integration regions for visual hyperacuity," *Vision Research*, 17, 89-93.
- Wilson, H.R., and Giese, S.C., [1977]
 "Threshold visibility of frequency grating patterns," *Vision Research*, 17, 1177-1190.
- Witkin, A.P., [1981]
 "Recovering surface shape and orientation from texture," *Artificial Intelligence*, 17, 17-45.
- Yosida, K., [1971]
Functional Analysis; third edition, Springer-Verlag, New York.
- Young, D.M., [1971]
Iterative Solution of Large Linear Systems, Academic Press, New York.
- Young, D.M., and Gregory, R.T., [1972, 1973]
A Survey of Numerical Mathematics, Vol. I, II, Addison-Wesley, Reading, MA.
- Zienkiewicz, O.C., [1974]
 "Constrained variational principles and penalty function methods in finite element analysis," *Conference on the Numerical Solution of Differential Equations*, Watson, G.A. (ed.), *Lecture Notes in Mathematics*, Vol. 363, Springer-Verlag, New York, 207-214.
- Zienkiewicz, O.C., [1977]
The Finite Element Method; Third edition, McGraw-Hill, London.
- Zucker, S.W., [1978]
 "Vertical and horizontal processes in low level vision," *Computer Vision Systems*, Riseman and Hanson (eds.), 187-195.
- Zucker, S.W., and Mohammed, J.L., [1979]
 "A hierarchical system for line labeling and grouping," *IEEE Conf. Image Proces. and Pat. Rec.*, 410-415.